

Brocade Fabric OS Command Reference, 8.0.1

Supporting Fabric OS 8.0.1

Copyright © 2016, Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, ClearLink, DCX, Fabric OS, HyperEdge, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, Vplane, and Vyatta are registered trademarks, and Fabric Vision is a trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of others.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Contents

Preface

| | |
|--|-------|
| Document conventions | .xv |
| Text formatting conventions | .xv |
| Command syntax conventions | .xv |
| Notes, cautions, and warnings | .xvi |
| Brocade resources | .xvi |
| Contacting Brocade Technical Support | .xvi |
| Brocade customers | .xvii |
| Brocade OEM customers | .xvii |
| Document feedback | .xvii |

About this Document

| | |
|--|-------|
| Supported hardware | xix |
| Brocade Gen 5 platform (16-Gbps) fixed-port switches | xix |
| Brocade Gen 5 platform (16-Gbps) DCX Directors | xix |
| Brocade Gen 6 platform (32-Gbps) fixed-port switches | .xx |
| Brocade Gen 6 platform (32-Gbps) Directors | .xx |
| What is new in this document | .xx |
| Changes made for the initial release (53-1004112-01) | .xx |
| Changes made for this release (53-1004112-02) | xxiii |
| CLI usage conventions | xxiv |

Using Fabric OS Commands

| | |
|---|---|
| Using the command line interface | 1 |
| Understanding Role-Based Access Control | 1 |
| Encryption commands and permissions | 2 |
| Understanding Virtual Fabric restrictions | 3 |
| Determining RBAC permissions for a specific command | 3 |

Fabric OS Commands

| | |
|----------------------------|----|
| aaaConfig | 5 |
| ad | 12 |
| ag | 22 |
| agAutoMapBalance | 45 |
| agShow | 48 |
| aliAdd | 51 |
| aliCreate | 52 |
| aliDelete | 53 |
| aliRemove | 54 |
| aliShow | 55 |
| appLoginHistory | 57 |

| | |
|------------------------------|-----|
| aptPolicy..... | 59 |
| auditCfg..... | 62 |
| auditDump..... | 64 |
| authUtil..... | 66 |
| bannerSet..... | 71 |
| bannerShow..... | 72 |
| bcastShow..... | 73 |
| bladeCfgGeMode..... | 74 |
| bladeDisable..... | 76 |
| bladeEnable..... | 78 |
| bladeSwap..... | 80 |
| bootLunCfg..... | 81 |
| bufOpMode..... | 83 |
| cfgActvShow..... | 87 |
| cfgAdd..... | 88 |
| cfgClear..... | 89 |
| cfgCreate..... | 90 |
| cfgDelete..... | 91 |
| cfgDisable..... | 92 |
| cfgEnable..... | 93 |
| cfgRemove..... | 95 |
| cfgSave..... | 96 |
| cfgShow..... | 98 |
| cfgSize..... | 102 |
| cfgTransAbort..... | 104 |
| cfgTransShow..... | 105 |
| chassisBeacon..... | 107 |
| chassisCfgPerrThreshold..... | 109 |
| chassisDisable..... | 111 |
| chassisDistribute..... | 112 |
| chassisEnable..... | 115 |
| chassisName..... | 116 |
| chassisShow..... | 117 |
| classConfig..... | 119 |
| cliHistory..... | 123 |
| configDefault..... | 125 |
| configDownload..... | 127 |

| | |
|-----------------------|-----|
| configList | 134 |
| configRemove | 135 |
| configShow | 136 |
| configUpload | 140 |
| configure | 145 |
| configureChassis..... | 157 |
| creditRecovMode | 161 |
| dataTypeShow | 166 |
| date | 167 |
| dbgShow | 169 |
| defZone | 170 |
| diagClearError | 172 |
| diagDisablePost..... | 173 |
| diagEnablePost | 174 |
| diagHelp..... | 175 |
| diagPost..... | 176 |
| diagRetry | 177 |
| diagShow | 178 |
| diagStatus | 179 |
| distribute | 180 |
| dlsReset..... | 182 |
| dlsSet | 183 |
| dlsShow | 188 |
| dnsConfig..... | 189 |
| enclosureShow | 191 |
| errClear | 193 |
| errDelimiterSet | 194 |
| errDump..... | 195 |
| errFilterSet..... | 200 |
| errModuleShow..... | 201 |
| errShow | 202 |
| ethIf | 206 |
| exit | 210 |
| extnCfg..... | 211 |
| fabRetryShow | 213 |
| fabRetryStats..... | 215 |
| fabricLog | 217 |

| | |
|---------------------------|-----|
| fabricName | 221 |
| fabricPrincipal | 223 |
| fabricShow | 226 |
| fabStatsShow | 229 |
| fanDisable | 231 |
| fanEnable | 232 |
| fanShow | 233 |
| faPwwn | 234 |
| fastBoot | 241 |
| fcipHelp | 243 |
| fcipLedTest | 244 |
| fcipPathTest | 245 |
| fcPing | 247 |
| fcLogClear | 255 |
| fcLogDisable | 256 |
| fcLogEnable | 257 |
| fcLogShow | 258 |
| fcProbeShow | 259 |
| fcRIsProbe | 261 |
| fcRIsShow | 263 |
| fcrBcastConfig | 265 |
| fcrConfigure | 267 |
| fcrEdgeShow | 270 |
| fcrFabricShow | 272 |
| fcrIcIPathBWMonitor | 274 |
| fcrLsan | 276 |
| fcrLsanCount | 279 |
| fcrLsanMatrix | 280 |
| fcrPhyDevShow | 286 |
| fcrProxyConfig | 288 |
| fcrProxyDevShow | 290 |
| fcrResourceShow | 292 |
| fcrRouterPortCost | 294 |
| fcrRouteShow | 296 |
| fcrXlateConfig | 298 |
| fddCfg | 302 |
| fdmiCacheShow | 304 |

| | |
|------------------------|-----|
| fdmiShow | 305 |
| femDump | 308 |
| ficonCfg | 310 |
| ficonClear | 312 |
| ficonCupSet | 313 |
| ficonCupShow | 316 |
| ficonHelp | 319 |
| ficonShow | 320 |
| fipsCfg | 329 |
| firmwareActivate | 334 |
| firmwareCheck | 335 |
| firmwareCleanInstall | 336 |
| firmwareCommit | 338 |
| firmwareDownload | 339 |
| firmwareDownloadStatus | 344 |
| firmwareKeyShow | 346 |
| firmwareRestore | 347 |
| firmwareShow | 348 |
| firmwareSync | 350 |
| flow | 351 |
| fosConfig | 373 |
| fosExec | 376 |
| frameLog | 380 |
| fspfShow | 385 |
| gePortErrShow | 388 |
| h | 389 |
| haDisable | 390 |
| haDump | 391 |
| haEnable | 393 |
| haFailover | 394 |
| haReboot | 395 |
| haRedundancy | 396 |
| haShow | 398 |
| haSyncStart | 400 |
| haSyncStop | 401 |
| help | 402 |
| historyLastShow | 404 |

| | |
|--------------------------|-----|
| historyMode | 406 |
| historyShow | 407 |
| i | 409 |
| iflShow | 411 |
| interfaceShow | 413 |
| iodReset | 420 |
| iodSet | 421 |
| iodShow | 422 |
| ipAddrSet | 423 |
| ipAddrShow | 430 |
| ipFilter | 434 |
| ipSecConfig | 440 |
| islShow | 459 |
| itemList | 462 |
| killTelnet | 464 |
| ldapCfg | 465 |
| lfCfg | 467 |
| licenseAdd | 472 |
| licenseIdShow | 473 |
| licensePort | 474 |
| licenseRemove | 477 |
| licenseShow | 479 |
| licenseSlotCfg | 482 |
| linkCost | 485 |
| logicalGroup | 488 |
| login | 492 |
| logout | 493 |
| lsanZoneShow | 494 |
| lsCfg | 497 |
| lsDbShow | 504 |
| mapsConfig | 508 |
| mapsDb | 513 |
| mapsHelp | 519 |
| mapsPolicy | 520 |
| mapsRule | 524 |
| mapsSam | 537 |
| memShow | 539 |

| | |
|------------------------------|-----|
| motd | 540 |
| msCapabilityShow | 541 |
| msConfigure | 542 |
| msPlatShow | 544 |
| msPlatShowDBC | 545 |
| msPIClearDB | 546 |
| msPIMgmtActivate | 547 |
| msPIMgmtDeactivate | 548 |
| msTdDisable | 549 |
| msTdEnable | 550 |
| msTdReadConfig | 551 |
| myId | 552 |
| nbrStateShow | 553 |
| nbrStatsClear | 556 |
| nodeFind | 558 |
| nsAliasShow | 560 |
| nsAllShow | 565 |
| nsCamShow | 566 |
| nsDevLog | 570 |
| nsShow | 573 |
| nsZoneMember | 578 |
| nsZoneShow | 582 |
| openSource | 584 |
| passwd | 585 |
| passwdCfg | 589 |
| pathInfo | 596 |
| pdShow | 604 |
| portAddress | 605 |
| portAlpaShow | 608 |
| portBeacon | 609 |
| portBufferCalc | 611 |
| portBufferShow | 613 |
| portCamShow | 616 |
| portCfg | 618 |
| portCfgAlpa | 651 |
| portCfgAutoDisable | 652 |
| portCfgCompress | 656 |

| | |
|--------------------------------|-----|
| portCfgCreditRecovery | 658 |
| portCfgDefault | 659 |
| portCfgDPort | 660 |
| portCfgEncrypt | 663 |
| portCfgEport | 665 |
| portCfgEportCredits | 667 |
| portCfgEXPort | 669 |
| portCfgFaultDelay | 673 |
| portCfgFec | 676 |
| portCfgFillword | 682 |
| portCfgFlogiLogout | 685 |
| portCfgFportBuffers | 686 |
| portCfgGE | 687 |
| portCfgGeMediaType | 690 |
| portCfgGport | 692 |
| portCfgISLMode | 693 |
| portCfgLongDistance | 695 |
| portCfgLossTov | 699 |
| portCfgLport | 701 |
| portCfgNonDfe | 703 |
| portCfgNPiVPort | 706 |
| portCfgNPort | 710 |
| portCfgOctetSpeedCombo | 712 |
| portCfgPersistence | 714 |
| portCfgPersistentDisable | 716 |
| portCfgPersistentEnable | 719 |
| portCfgQoS | 722 |
| portCfgShow | 725 |
| portCfgSpeed | 739 |
| portCfgTdz | 742 |
| portCfgTrunkPort | 744 |
| portCfgVEXPort | 745 |
| portChannelShow | 748 |
| portCmd | 750 |
| portDebug | 762 |
| portDecom | 763 |
| portDisable | 764 |

| | |
|--------------------------|-----|
| portDPortTest | 767 |
| portEnable | 776 |
| portEncCompShow | 778 |
| portErrShow | 780 |
| portFlagsShow | 783 |
| portLedTest | 785 |
| portLogClear | 787 |
| portLogConfigShow | 788 |
| portLogDisable | 789 |
| portLogDump | 790 |
| portLogDumpPort | 792 |
| portLogEnable | 793 |
| portLogEventShow | 794 |
| portLoginShow | 795 |
| portLogPdisc | 797 |
| portLogReset | 798 |
| portLogResize | 799 |
| portLogShow | 800 |
| portLogShowPort | 807 |
| portLogTypeDisable | 808 |
| portLogTypeEnable | 809 |
| portLoopbackTest | 810 |
| portName | 814 |
| portPeerBeacon | 817 |
| portPerfShow | 820 |
| portRouteShow | 824 |
| portShow | 826 |
| portStats64Show | 889 |
| portStatsClear | 894 |
| portStatsShow | 896 |
| portSwap | 904 |
| portSwapDisable | 906 |
| portSwapEnable | 907 |
| portSwapShow | 908 |
| portTest | 909 |
| portTestShow | 913 |
| portTrunkArea | 916 |

| | |
|-----------------------------|------|
| portZoneShow | 922 |
| powerOffListSet | 923 |
| powerOffListShow | 925 |
| psShow | 927 |
| rasAdmin | 928 |
| rasMan | 931 |
| reboot | 933 |
| relayConfig | 935 |
| roleConfig | 937 |
| rootAccess | 940 |
| routeHelp | 942 |
| rtLogTrace | 943 |
| sddQuarantine | 944 |
| secActiveSize | 946 |
| secAuthSecret | 947 |
| secCertUtil | 949 |
| secCryptoCfg | 960 |
| secDefineSize | 966 |
| secGlobalShow | 967 |
| secHelp | 969 |
| secPolicyAbort | 970 |
| secPolicyActivate | 971 |
| secPolicyAdd | 972 |
| secPolicyCreate | 975 |
| secPolicyDelete | 979 |
| secPolicyDump | 981 |
| secPolicyFCSMove | 984 |
| secPolicyRemove | 986 |
| secPolicySave | 988 |
| secPolicyShow | 989 |
| secStatsReset | 991 |
| secStatsShow | 993 |
| sensorShow | 995 |
| serDesTuneMode | 996 |
| setContext | 999 |
| setDbg | 1000 |
| setVerbose | 1001 |

| | |
|----------------------------|------|
| sfpShow | 1002 |
| shellFlowControlDisable | 1013 |
| shellFlowControlEnable | 1014 |
| slotPowerOff | 1015 |
| slotPowerOn | 1016 |
| slotShow | 1017 |
| snmpConfig | 1022 |
| snmpTraps | 1036 |
| spinFab | 1039 |
| sshUtil | 1043 |
| statsClear | 1049 |
| stopPortTest | 1050 |
| supportFfdc | 1051 |
| supportFtp | 1052 |
| supportInfoClear | 1054 |
| supportSave | 1055 |
| supportShow | 1059 |
| supportShowCfgDisable | 1062 |
| supportShowCfgEnable | 1064 |
| supportShowCfgShow | 1066 |
| switchBeacon | 1068 |
| switchCfgPersistentDisable | 1070 |
| switchCfgPersistentEnable | 1071 |
| switchCfgSpeed | 1072 |
| switchCfgTrunk | 1074 |
| switchDisable | 1075 |
| switchEnable | 1076 |
| switchName | 1077 |
| switchShow | 1079 |
| switchUptime | 1097 |
| switchViolation | 1098 |
| syslogAdmin | 1099 |
| sysShutDown | 1102 |
| tempShow | 1104 |
| timeOut | 1106 |
| topologyShow | 1107 |
| traceDump | 1111 |

| | |
|-------------------|------|
| trunkDebug | 1114 |
| trunkShow | 1115 |
| tsClockServer | 1119 |
| tsTimeZone | 1121 |
| turboRamTest | 1125 |
| upTime | 1127 |
| uRouteShow | 1128 |
| usbStorage | 1130 |
| userConfig | 1132 |
| version | 1139 |
| wwn | 1140 |
| wwnAddress | 1142 |
| wwnRecover | 1145 |
| zone | 1147 |
| zoneAdd | 1159 |
| zoneCreate | 1161 |
| zoneDelete | 1164 |
| zoneHelp | 1165 |
| zoneObjectCopy | 1166 |
| zoneObjectExpunge | 1167 |
| zoneObjectRename | 1169 |
| zoneObjectReplace | 1170 |
| zoneRemove | 1172 |
| zoneShow | 1174 |

Primary FCS Commands

| | |
|----------------------|------|
| Primary FCS commands | 1181 |
|----------------------|------|

Command Availability

| | |
|--|------|
| Command validation checks | 1183 |
| Command Virtual Fabrics and Remote Domain execution restrictions | 1184 |

Preface

- [Document conventions](#) xv
- [Brocade resources](#) xvi
- [Contacting Brocade Technical Support](#) xvi
- [Document feedback](#) xvii

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Brocade technical documentation.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used in the flow of the text to highlight specific words or phrases.

| Format | Description |
|--------------|---|
| bold text | Identifies command names |
| | Identifies keywords and operands |
| | Identifies the names of user-manipulated GUI elements |
| | Identifies text to enter at the GUI |
| italic text | Identifies emphasis |
| | Identifies variables |
| | Identifies document titles |
| Courier font | Identifies CLI output |
| | Identifies command syntax examples |

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

| Convention | Description |
|-------------|--|
| bold text | Identifies command names, keywords, and command options. |
| italic text | Identifies a variable. |
| value | In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, --show WWN. |
| [] | Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets. |

| | |
|---------------|---|
| { x y z } | A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options. In Fibre Channel products, square brackets may be used instead for this purpose. |
| x y | A vertical bar separates mutually exclusive elements. |
| < > | Nonprinting characters, for example, passwords, are enclosed in angle brackets. |
| ... | Repeat the previous element, for example, member[member...]. |
| \ | Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash. |

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE

A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.



CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Brocade resources

Visit the Brocade website to locate related documentation for your product and additional Brocade resources.

You can download additional publications supporting your product at www.brocade.com. Select the Brocade Products tab to locate your product, then click the Brocade product name or image to open the individual product page. The user manuals are available in the resources module at the bottom of the page under the Documentation category.

To get up-to-the-minute information on Brocade products and resources, go to MyBrocade. You can register at no cost to obtain a user ID and password.

Release notes are available on MyBrocade under Product Downloads.

White papers, online demonstrations, and data sheets are available through the Brocade website.

Contacting Brocade Technical Support

As a Brocade customer, you can contact Brocade Technical Support 24x7 online, by telephone, or by e-mail. Brocade OEM customers contact their OEM/Solutions provider.

Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to <http://www.brocade.com/services-support/index.html>.

If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

| Online | Telephone | E-mail |
|--|---|---|
| Preferred method of contact for nonurgent issues: <ul style="list-style-type: none"> • My Cases through MyBrocade • Software downloads and licensing tools • Knowledge Base | Required for Sev 1-Critical and Sev 2-High issues: <ul style="list-style-type: none"> • Continental US: 1-800-752-8061 • Europe, Middle East, Africa, and Asia Pacific: +800-AT FIBREE (+800 28 34 27 33) • For areas unable to access toll free number: +1-408-333-6061 • Toll-free numbers are available in many countries. | support@brocade.com Please include: <ul style="list-style-type: none"> • Problem summary • Serial number • Installation details • Environment description |

Brocade OEM customers

If you have purchased Brocade product support from a Brocade OEM/Solution Provider, contact your OEM/Solution Provider for all of your product support needs.

- OEM/Solution Providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/Solution Provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/Solution Provider.

Document feedback

To send feedback and report errors in the documentation you can use the feedback form posted with the document or you can e-mail the documentation team.

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. You can provide feedback in two ways:

- Through the online feedback form in the HTML documents posted on www.brocade.com.
- By sending your feedback to documentation@brocade.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

About this Document

- [Supported hardware](#) xix
- [What is new in this document](#) xx
- [CLI usage conventions](#) xxiv

Supported hardware

In those instances in which procedures or parts of procedures documented here apply to some devices but not to others, this list identifies exactly which devices are supported by Fabric OS 8.0.1 and which are not

Although many different software and hardware configurations are tested and supported by Brocade Communication Systems, Inc for Fabric OS 8.0.1, documenting all possible configurations and scenarios is beyond the scope of this document.

Fabric OS 8.0.1 provides features that interoperate with the Brocade Analytics Monitoring Platform running Fabric OS 7.4.0_amp. For more information about the Brocade Analytics Monitoring Platform and Fabric OS 8.0.1, you can find content on MyBrocade.com.

The following hardware platforms are supported by Fabric OS 8.0.1.

Brocade Gen 5 platform (16-Gbps) fixed-port switches

- Brocade 6505 switch
- Brocade 6510 switch
- Brocade 6520 switch
- Brocade M6505 blade server SAN I/O module
- Brocade 6543 blade server SAN I/O module
- Brocade 6545 blade server SAN I/O module
- Brocade 6546 blade server SAN I/O module
- Brocade 6547 blade server SAN I/O module
- Brocade 6548 blade server SAN I/O module
- Brocade 6558 blade server SAN I/O module
- Brocade 7840 Extension Switch

Brocade Gen 5 platform (16-Gbps) DCX Directors

NOTE

For ease of reference, Brocade chassis-based storage systems are standardizing on the term "Director". The legacy term "Backbone" can be used interchangeably with the term "Director".

- Brocade DCX 8510-4 Director
- Brocade DCX 8510-8 Director

What is new in this document

Brocade Gen 6 platform (32-Gbps) fixed-port switches

- Brocade G620

Brocade Gen 6 platform (32-Gbps) Directors

- Brocade X6-4 Director
- Brocade X6-8 Director

What is new in this document

The Virtual Fabric restrictions for Fabric OS commands table in the “[Command Availability](#)” has been updated. The table lists the Context, Switch type, and Remote Domain Execution for the commands.

NOTE

Admin Domain is deprecated and the information about Admin Domain restrictions are removed.

Changes made for the initial release (53-1004112-01)

New commands

- **diagStatus** - Displays currently running diagnostic tests.
- **haReboot** - Performs high availability (HA) reboot.
- **openSource** - Displays open-source licenses.
- **rootAccess** - Configures root access through CLI.

Modified commands

- **aaaConfig** - Added a note that the **-e** operand is applicable only to RADIUS configuration.
- **ad** - The following error message is displayed when you execute this command:
Warning: Admin Domains are not supported in Fabric OS v8.0.1.
Admin Domain commands and functionality will be removed in future Fabric OS versions.
- **authUtil** - Added a note that before setting the authentication type using the **-a** operand, you must make sure that FCAP certificates and dhchap secrets are configured. The **--authinit** command does not initiate authentication on disabled ports.
- **bladeCfgGeMode** - This command is supported only on the Brocade FX8-24 blade in a Brocade DCX 8510 chassis.
- **cfgSize** - Removed all references to the term 'Admin Domain'.
- **chassisBeacon** - Added a note that on a Brocade X6 Director, enabling chassis beaconing activates beaconing on CP blades also.
- **chassisName** - Added a note that the chassis name configured using this command appears in all the chassis event-related RASLog messages.
- **configDownload** - Added the steps to perform to backup and then restore a configuration in a switch using Virtual Fabrics.
- **configUpload** - Added the steps to perform to backup and then restore a configuration in a switch using Virtual Fabrics.
- **configure** - Added “F-Port Device Update Mode” field.
- **diagShow** - Added a note that the **-post** operand is supported only on 32 Gbps-capable platforms.
- **distribute** - Added a note that if FCS policy is enabled, only primary FCS switch can distribute the data.

- **extnCfg** - Added the `--config -default` and `--config -clear` operands.
- **fanDisable** - Added note that the fan units on a Brocade X6 Director do not change their RPM speed to 0 and they take few minutes to change the speed to minimum RPM after being disabled.
- **faPwwn** - Added a note that the ports must be disabled before executing the `--enable` and `--assign` options.
- **fcipLedTest** - Added the behavior of this command on the Brocade SX6 extension blade.
- **fcipPathTest** - Added details for Brocade SX6 extension blade.
- **fcrEdgeShow** - This command displays information only for online EX_Ports.
- **fddCfg** - Added a note that when a policy is set to strict mode, the manual distribution of the policy is not allowed.
- **firmwareDownload** - Removed the following unsupported operands: `-a`, `-t`, `-e`, and `-o`.
- **flow** - Added the predefined frame type 'Srr'.
- **fosExec** - Added a note that the `--domain` operand can be executed only if the remote Fosexec feature is enabled on local and remote switches.
- **haShow** - Added slot number details for the Brocade X6 Director. Added **haShow** command example on the Brocade X6 Director.
- **mapsConfig** - Removed **sw_critical** and **sw_marginal** MAPS actions. These actions are always enabled by default.
- **mapsDb** - Added new fields 'fenced circuits' and 'ports with highest zoned device ratio' fields.
- **mapsRule** - Added the following monitoring systems under the respective category"
 - **GE Port Health**: GE_CRC, GE_INV_LEN, GE_LOS_OF_SIG, STATE_CHG, UTIL, and IP_UTIL.
 - **FCIP health**: IP_UTIL, IP_PKTLOSS, IP_RTT, and IP_JITTER.
 - **Fabric performance impact**: CRED_ZERO, THPUT_DEGRADE, and IT_FLOW.
 - **Security health**: DAYS_TO_EXPIRE, EXPIRED_CERTS, DID_CHG, FLOGI, and FAB_CFG.
 - **Fabric state change**: TEMP, ETH_MGMT_PORT_STATE, and FLASH_USAGE.
 - **Switch status policy**: FAN_AIRFLOW_MISMATCH.
- **msConfigure** - Removed the following note: This command is supported only in ADO and AD255 contexts.
- **msPIClearDB** - Removed the following note: This command is supported only in ADO and AD255 contexts.
- **msPIMgmtActivate** - Removed the following note: This command is supported only in ADO and AD255 contexts.
- **msPIMgmtDeactivate** - Removed the following note: This command is supported only in ADO and AD255 contexts.
- **msTdDisable** - Removed the following note: This command is supported only in ADO and AD255 contexts.
- **msTdEnable** - Removed the following note: This command is supported only in ADO and AD255 contexts.
- **nodeFind** - The command is enhanced to accept WWN as single hexadecimal value.
- **nsShow** - Added the 'Device link speed' field.
- **passwdCfg** - Added the `-charset`, `-allowuser`, `--hash`, and `--showhash` options.
- **portCfg** - Added details for Brocade SX6 blade and removed all references to Brocade 7800. Added the following options: **sla**, **--proto-app**, **--rst-propagation**, **--non-terminated**, and **--load-leveling**. Removed the **--legacy** option.
- **portCfgCompress** - Added the following note: On 16 Gbps-capable platforms, both encryption and compression features are supported. On 32 Gbps-capable platforms, only compression is supported and encryption is not supported.
- **portCfgDefault** - Added the `--help` option.
- **portCfgEncrypt** - Added the following note: On 16 Gbps-capable platforms, both encryption and compression features are supported. On 32 Gbps-capable platforms, only compression is supported and encryption is not supported.

What is new in this document

- **portCfgEportCredits** - Modified the description of the credits option. Modified the description to state that E_Port credits are mutually exclusive with Buffer Optimized mode. ICL ports are supported with this command. Added valid credit range for ICL ports.
- **portCfgGE** - Added details for Brocade SX6 blade and removed Brocade 7800.
- **portCfgNonDfe** - Updated the description and the **--show** output. Added new '8G Non-DFE State' field.
- **portCfgQoS** - The **--setratelimit** operand is not supported on Brocade X6-4 Director and Brocade X6-8 Director.
- **portCfgShow** - Added details for Brocade SX6 blades and removed all references to Brocade 7800. Added the following options: **lag**, **tcl**, and **sla**.
- **portCfgSpeed** - Added 32 Gbps fixed speed. Added a note that speed 'ax' is not supported on 16G and 32G platforms and 's' is not supported on 32G platforms.
- **portCmd** - Added support for Brocade SX6 extension blade. Added the following options: **--uni-directional**, **--dscp**, **--l2cos**, **--summary**, **-connection**, **-detail**, **-historical**, and **--all-ls**.
- **portDPortTest** - Added **--set [ewrap | owrap]** and **--clear [ewrap | owrap]** options. Added the following note: The D_Port long duration test can only be run on a single port at a time.
- **portEncCompShow** - Removed the 'Active Speed' field from the output. Added the following note: On 16 Gbps-capable platforms, both encryption and compression features are supported. On 32 Gbps-capable platforms, only compression is supported and encryption is not supported.
- **portLedTest** - Removed the **ceeporledtest** command.
- **portLoopbackTest** - Removed the **ceeporloopbacktest** command.
- **portName** - Added **-d** option.
- **portShow** - Added details for Brocade SX6 blade and removed Brocade 7800. Added the following options: **sla**, **--link-local**, **--sort**, **-reset**, **-lifetime**, **--hist-stats**, and **--filter**.
- **portStats64Show** - The **stat64_fec_cor** field is displayed only on 16 Gbps-capable platforms and not on 32 Gbps-capable platform.
- **portStatsShow** - Added **fec_corrected_rate** field. The **fec_cor_detected** and **tim_latency_vc** fields are displayed only on 16 Gbps-capable platforms and not on 32 Gbps-capable platform. The following fields are displayed only on 32 Gbps platforms: **lat_tot_pkt_vc**, **lat_hi_time_vc**, **lat_lo_time_vc**, and **latency_dma_ts**.
- **portTest** - This command is supported on the N->N Loopback Ports and it is not supported on the F_Ports on VF-enabled switch.
- **powerOffListSet** - Updated the command example.
- **powerOffListShow** - Updated the command example.
- **reboot** - Added details for Brocade X6 Director.
- **secCertUtil** - Added the following options: **-hash**, **-https**, **-type**, **-years**, **-nowarn**.
- **secCryptoCfg** - Added new options to manage template files.
- **sfpShow** - Corrected command example for QSFP.
- **slotShow** - Added new blade IDs: 175, 177, 178, and 186. Removed unsupported blade IDs.
- **snmpConfig** - This command has been enhanced to run with command line options (non-interactive).
- **spinFab** - This command is supported on AE_Ports and not supported on core blade connected to port blade.
- **sshUtil** - Added the following options: **genhostkey**, **showhostkey**, and **delethehostkey**. Added **-all** option for the delknownhost operand.
- **switchCfgSpeed** - Added 32 Gbps speed.
- **switchName** - The command is enhanced to accept switch name beginning with number.

- **switchShow** - Added the following disabled reasons: Disabled (EX_Port IR POD License Limit Exceeded), Disabled (License not Installed for Integrated Routing), and Disabled (Incompatible AMP Version). Added 'ESC mismatch, AMP Version Conflict' and 'Incompatible AMP Version' segmentation reasons.
- **traceDump** - Added the 'Reboot' tracedump type.
- **turboRamTest** - Removed the **ceeturboramtest** command.
- **userConfig** - Added **-at** | **-access** option. An asterisk (*) next to the account name in the **--show** and **-a** output indicates that the password for that user account is still at the default value.

Deprecated commands

- **bpPortLoopbackTest**
- **bpTurboRamTest**
- **ceePortLedTest**
- **ceePortLoopbackTest**
- **ceeTurboRamTest**
- **cmsh**
- **fcoe**
- **fcoeLoginCfg**
- **fcoeLoginGroup**
- **iclCfg**
- **ifModeSet**
- **ifModeShow**
- **portMirror**

Changes made for this release (53-1004112-02)

Modified commands

- **ad** - Added a note that Admin Domains are not supported in Fabric OS v8.0.1. When this command is executed, it displays the following warning message:
Warning: Admin Domains are not supported in Fabric OS v8.0.1.
Admin Domain commands and functionality will be removed in future Fabric OS versions.
- **bufOpMode** - Added details for Brocade SX6 blade.
- **errDump** - Updated command examples for the **--start** and **--stop** options.
- **errShow** - Updated command examples for the **--start** and **--stop** options.
- **extnCfg** - Added **-all** options to specify all slots in the chassis.
- **haReboot** - Added the following note: When the switch-based system is not fully up or stable, the **haReboot** command cannot be completed.
- **portCfgDPort** - D_Port configuration is not supported on SIM-Ports.
- **portCfgGeMediaType** - Added details for Brocade 7840 and removed Brocade 7800.
- **portCfgNonDfe** - Made editorial changes.

- **switchName** - Added a note in the description of the *name* field that the switchname which begins with numeric character should have atleast one alphabet. Switchname with all numeric character is not valid. Example switchname 12345678 is invalid.
- **zone** - Removed all references of Admin Domain.

CLI usage conventions

- Some Fabric OS show commands display certain switch or fabric components in three-digit decimal numbers, for example, 003. When you use these numeric identifiers as inputs to other commands, you must remove the leading zeros; otherwise the commands will either fail or generate incorrect results. Adding zeros to any decimal number in a command input will cause that number to be treated as an Octal number.
- The entire command line (both commands and options) is case-sensitive. Selected command names and options may also support Java-style capitalization. Java-style capitalization means that while bannershow and bannerShow will both work, BANNERSHOW and BannerShow will not. The options of commands on the other hand, are strictly case sensitive and therefore only lower case or Java-style capitalization can be entered depending on each option. This means that while the -USB option of the configDownload command works, but -usb will not. Refer to each syntax for explicit instructions on supported capitalization for each command and options.
- When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required. Example: zonecreate "zonename" requires that the value for zonename be in quotation marks.
- Automatic page breaks in CLI command output displays are being phased out. Use the more option to display command output with page breaks: command | more. Do not use the more option in conjunction with help pages. Executing help command | more will display a "no manual entry for command" message.

Using Fabric OS Commands

- [Using the command line interface](#) 1
- [Understanding Role-Based Access Control](#)..... 1
- [Understanding Virtual Fabric restrictions](#)..... 3
- [Determining RBAC permissions for a specific command](#) 3

Using the command line interface

The Fabric OS command line interface (CLI), accessed via Telnet, SSH, or serial console, provides full management capability on a Brocade switch. The Fabric OS CLI enables an administrator to monitor and manage individual switches, ports, and entire fabrics from a standard workstation. Selected commands must be issued from a secure Telnet or SSH session.

Access is controlled by a switch-level password for each access level. The commands available through the CLI are based on the user's login role and the license keys used to unlock certain features.

The Fabric OS CLI provides the following capabilities:

- Access to the full range of Fabric OS features, given the license keys installed.
- Assistance with configuration, monitoring, dynamic provisioning, and daily management of every aspect of storage area networks (SANs).
- A deeper view of the tasks involved in managing a Brocade SAN.
- Identification, isolation, and management of SAN events across every switch in the fabric.
- Management of Brocade licenses.

The documentation for each command includes a synopsis of its syntax, a description of command use, and a set of examples. The same information can be accessed by issuing the **help** command followed by the command name on a Brocade switch or director. This command displays the help page for the specified command. For example, to display the help page for **portCfg**, enter:

```
switch:admin> help portCfg
```

Understanding Role-Based Access Control

Fabric OS implements Role-Based Access Control (RBAC) to control access to all Fabric OS operations.

Eight predefined roles are supported, as described in [Table 1](#). These predefined role definitions are guided by perceived common operational situations and the operations and effects a role is permitted to have on a fabric and individual fabric elements.

TABLE 1 Role definitions

| Role name | Definition |
|-------------|---|
| User | Non-administrative use, such as monitoring system activity. In Fabric OS v6.2.0 and later, the user account gains access to Fabric ID 128. This is the default logical fabric after a firmware upgrade. |
| Operator | A subset of administrative tasks typically required for routine maintenance operations. |
| SwitchAdmin | Administrative use excluding security, user management, and zoning. |
| ZoneAdmin | Zone management only. |

TABLE 1 Role definitions (Continued)

| Role name | Definition |
|------------------|--|
| FabricAdmin | Administrative use excluding user management management. |
| BasicSwitchAdmin | A subset of administrative tasks, typically of a more limited scope and effect. |
| Admin | All administrative tasks, including encryption and chassis commands. |
| SecurityAdmin | Administrative use including admin, encryption, security, user management, and zoning. |

In addition to these predefined roles, Fabric OS v7.0.0 and later provides support for creating user-defined roles. Refer to the **roleConfig** command for more information.

Additional command restrictions apply depending on whether Virtual Fabrics are enabled in a fabric. Refer to ["Command Availability"](#).

Encryption commands and permissions

There are two system RBAC roles that are permitted to perform encryption operations.

- Admin and SecurityAdmin
Users authenticated with the Admin and SecurityAdmin RBAC roles may perform cryptographic functions assigned to the FIPS Crypto Officer, including the following:
 - Perform encryption node initialization.
 - Enable cryptographic operations.
 - Manage critical security parameters (CSPs) input and output functions.
 - Zeroize encryption CSPs.
 - Register and configure a key vault.
 - Configure a recovery share policy.
 - Create and register recovery share.
 - Encryption group and clustering-related operations.
 - Manage keys, including creation, recovery, and archiving functions.
- Admin and FabricAdmin
Users authenticated with the Admin and FabricAdmin RBAC roles may perform routine encryption switch management functions including the following:
 - Configure virtual devices and crypto LUN.
 - Configure LUN/tape associations.
 - Perform re-keying operations.
 - Perform firmware download.
 - Perform regular Fabric OS management functions.

For a listing of RBAC permissions for cryptoCfg subcommands, refer to the *Fabric OS Encryption Administrator's Guide*.

Understanding Virtual Fabric restrictions

All Fabric OS commands are subject to additional RBAC enforcement with regard to Virtual Fabric contexts and switch types. Commands can be executed in one or more of the contexts described in [Table 2](#). Execution of chassis commands requires chassis permissions.

TABLE 2 Virtual Fabric contexts

| Context type | Definition |
|----------------------------|---|
| Switch context | Command applies to the current logical switch only, or to a specified logical switch. |
| Chassis context | Command applies to the chassis on which it is executed. |
| Switch and chassis context | Command can be executed in a logical switch context or in a chassis context. |
| Disallowed | Command is not supported in Virtual Fabric mode. |

Switch commands are further defined by the switch type restrictions as described in [Table 3](#). Switch type restrictions are not applicable to commands that require chassis permissions.

TABLE 3 Switch types

| Switch type | Definition |
|---------------------|---|
| All Switches | Command can be executed in any switch context. |
| Base Switch Only | Command can be executed only on the base switch. |
| Default Switch Only | Command can be executed only on the default switch. |
| N/A | Command is a chassis command or not supported in Virtual Fabric mode. |

In a Virtual Fabric environment where contexts are enforced, the following Virtual Fabric restrictions apply to the RBAC permissions specified in [Table 1](#). Refer to the **userConfig** command for more information on configuring user account access permissions in a Virtual Fabric environment.

- Any given role is allowed to execute all switch commands to which the role is authorized in the account's home context. The default home context is the default logical fabric FID 128.
- You can change an account's home context to a specified FID and configure the account permissions to access additional logical switches specified in the user's Fabric ID list.
- Accounts with user or admin permissions can be granted chassis permissions. A user account with the chassis role can execute chassis-level commands at the user RBAC access level. An admin account with the chassis role can execute chassis-level commands at the admin RBAC access level.

Use the **classConfig --showcli** command to look up the Virtual Fabric context for a specified command. Refer to ["Command Availability"](#) for a complete listing of Virtual Fabric restrictions that apply to the commands included in this manual.

Determining RBAC permissions for a specific command

To determine RBAC permission for a specific command, use the **classconfig** command.

1. Enter the **classconfig --showcli** command for a specified command.

The command displays the RBAC class and access permissions for each of the command options. Note that options for a single command option can belong to different classes.

2. Enter the **classconfig --showroles** command and specify the RBAC class of the command option you want to look up.

The command displays the default roles and the permissions they have to access commands in the specified RBAC class.

Determining RBAC permissions for a specific command

The following example shows how you can obtain permission information for the **zone** command. Suppose you want to know if a user with the SwitchAdmin role can create a zone. You issue the **classconfig --showcli** command for the zone command, which shows that the **zone --add** command belongs to the RBAC class "zoning". You then issue the **classconfig --showroles** command for the zoning RBAC class. The output shows that the SwitchAdmin role has 'Observe' (O) permissions only for any command in the zoning class. This means that the user with the SwitchAdmin role is not allowed to create zones. To allow this user to create a zone, you must change the user's access to any of the roles that have "observe and modify" (OM) access. Use the **userConfig** command to change the user's role or use the **roleConfig** command to create a custom role.

```
switch:admin> classconfig --showcli zone
CLI      Option          Permission  RBAC Class  Context
-----
zone     Killall         OM          Debug       vf
zone     evlogclear     OM          Debug       vf
zone     evlogshow      O           Debug       vf
zone     evlogtoggle    OM          Debug       vf
zone     mergeshow      O           Debug       vf
zone     stateshow      O           Debug       vf
zone     activate       OM          Zoning      vf
zone     add            OM          Zoning      vf
zone     copy           OM          Zoning      vf
zone     create         OM          Zoning      vf
zone     deactivate     OM          Zoning      vf
(output truncated)
```

```
switch:admin> classconfig --showroles zoning
Roles that have access to the RBAC Class 'zoning' are:
```

| Role Name | Permission |
|------------------|------------|
| ----- | ----- |
| User | O |
| Admin | OM |
| Factory | OM |
| Root | OM |
| Operator | O |
| SwitchAdmin | O |
| ZoneAdmin | OM |
| FabricAdmin | OM |
| BasicSwitchAdmin | O |
| SecurityAdmin | O |

NOTE

If a Role Name does not appear in the list, it indicates that the role is not available to the specified class and associated commands in that class are restricted and cannot be executed in that role.

Fabric OS Commands

aaaConfig

Manages RADIUS, LDAP, and TACACS+ configuration information.

Synopsis

aaaconfig

aaaconfig --show

aaaconfig --add | --change *server* -conf radius | ldap | tacacs+ [-p *port*] [-d *domain*] [-t *timeout*] [-s *secret*] [-a chap | pap | peap-mschapv2] [-e *encr_type*]

aaaconfig --remove *server* -conf radius | ldap | tacacs+

aaaconfig --move *server* -conf radius | ldap | tacacs+ *to_position*

aaaconfig --authspec *aaa1*[:*aaa2*] [-backup] [-nologout]

aaaconfig --help

Description

Use this command to manage the RADIUS, LDAP, and TACACS+ server configuration for the authentication, authorization and accounting (AAA) services. Use this command to display, add, remove, change, enable or disable the RADIUS, LDAP, or TACACS+ configuration.

Brocade switches use a local as well as a remote authentication mechanism for validating a login. Supported authentication protocols include Password Authentication Protocol (PAP), Challenge-Handshake Authentication Protocol (CHAP) and Protected Extensible Authentication Protocol (PEAP). Light-weight Directory Access Protocol (LDAP) authentication against Active Directory for user authentication and authorization is supported.

RADIUS, LDAP or TACACS+ servers are contacted in the order they appear in the configuration list. The first server returning authentication success or failure causes the authentication request to succeed or fail. If no response is received within the specified timeout, the next RADIUS, LDAP, or TACACS+ server in the list is contacted. An event entry logs if all RADIUS, LDAP, or TACACS+ servers fail to respond.

When the command succeeds, it triggers an event log (the Fabric OS error log) to indicate a server is added, removed, or modified. Refer to the *Fabric OS Message Reference* manual for specific details.

There are two modes of operation in LDAP authentication, FIPS mode and non-FIPS mode. However, there is no option to configure LDAP while the switch is in FIPS mode. The LDAP client checks if FIPS mode is set on the switch and uses FIPS-compliant TLS ciphers for LDAP. If FIPS mode is not set and the ADir server is configured for FIPS ciphers, it uses FIPS-compliant ciphers.

Configuration changes are persistently saved and take effect with the next AAA request. The configuration applies to all switch instances in a platform supporting multiple switch domains.

Notes Customers can use centralized RADIUS servers to manage AAA services for a switch, as defined in the RFC 2865 RADIUS specification.

This command can be executed when logged in through the console, Telnet or SSH connection.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

server

Specifies an IP address or a server name in dot-decimal notation. IPv6 addresses are supported. If a name is used, a DNS entry must be correctly configured for the server. If the specified server IP address or name already exists in the current configuration, the command fails and generates an error. However, the command does not validate the server name against the IP address in the configuration. Make sure to avoid duplicate configuration of the same server, one specified by the name, the other specified by the IP address.

--show

Displays the current AAA service configuration.

--add | --change *server* [*options*]

Adds or modifies a RADIUS, LDAP, or TACACS+ server. The **--add** option appends the specified server to the end of the current configuration list. A maximum of 5 servers are supported for each authentication type. The **--change** option modifies the specified server configuration to use the new arguments. The server must be one of the IP addresses or names shown in the current configuration.

The following *options* are supported:

-conf *radius* | *ldap* | *tacacs+*

Specifies the server configuration as either RADIUS, LDAP, or TACACS+. This operand is required.

The following operands are optional:

-p *port*

Specifies the RADIUS, LDAP, or TACACS+ server port number. Supported range is 1 to 65535. The default port is 1812 for RADIUS authentication. The default port is 1813 for RADIUS accounting. The default port is 389 for LDAP authentication. The default port is 49 for TACACS+ authentication. This operand is optional. If no port is specified, the default is used.

-t *timeout*

Specifies the response timeout for the RADIUS, the LDAP or the TACACS+ server. The supported range is 1 to 30 seconds. The default is 3 seconds. This operand is optional. If no timeout is specified, the default is used.

-d *domain*

Specifies the Windows domain name for the LDAP server, for example, brocade.com. This option is valid only with the **-conf *ldap*** option. This operand is required.

- s *secret*** Specifies a common secret between the switch and the RADIUS or TACACS+ server. The secret must be 8 to 40 characters for RADIUS server, and 1 to 40 characters for TACACS+ server. This option is valid only with **-conf radius** or **-conf tacacs+** options, and it is optional. The default value is **sharedsecret**.
- a** Specifies the remote authentication protocol for the RADIUS or TACACS+ server. This operand is valid with **-conf radius** or **-conf tacacs+** options, and it is optional. The default value for this operand is **CHAP**.
- Note that the distinction between protocols is only applicable to the packets between a system and the RADIUS or TACACS+ server. To authenticate a user to the system, a password is always used.
- Valid protocols are one of the following:
- pap** Password Authentication Protocol
- chap** Challenge-Handshake Authentication Protocol
- peap-mschapv2** Protected Extensible Authentication Protocol. This is applicable only to RADIUS configuration. The combination of "peap-mschapv2" and IPv6 causes RADIUS authentication to be rejected. PEAP with IPv4 succeeds.
- e *encr_type*** Specifies the encryption algorithm. This is applicable only to RADIUS configuration. Valid algorithms include the following:
- none** Null encryption algorithm
- aes256** AES-256 algorithm
- remove *server*** Removes the specified server from the configuration. The server must match one of the IP addresses or the names shown in the current configuration. The following operand is required:
- conf radius | ldap | tacacs+** Specifies the server configuration as either RADIUS, LDAP, or TACACS+. If the server is enabled, the command does not allow the last server to be removed from the configuration list. RADIUS, LDAP, or TACACS+ must first be disabled before the last server of the specified type may be removed.

--move server option

Moves the specified server from the current position in a RADIUS, LDAP, or TACACS+ configuration list to the specified position. If the specified position is the same as the current position, no change takes place. Valid *options* include the following:

-conf radius | ldap | tacacs+

Specifies the server configuration as RADIUS, LDAP, or TACACS+. This operand is required.

to_position

Specifies the new position for the server. The value for to_position is an integer, and must be within the range of server positions in the current configuration. Use the **--show** option to determine current server positions. This operand is required.

--authspec "aaa1[:aaa2]" [-backup] [-nologout]

Replaces the configuration with the specified AAA service. Each service can be specified only once in the list, for example, "radius; local; radius" is invalid. No edit option is provided. The **--authspec** option takes as an argument a semicolon-separated list of AAA services. Services must be enclosed in double quotation marks.

The following AAA services and service pairs are valid:

"local"

Default setting. Authenticates the user against the local database only. If the password does not match or the user is not defined, the login fails.

"radius"

When "radius" is specified, the first RADIUS server is contacted. If the RADIUS server is not reachable, the next RADIUS server is contacted. If the authentication fails, the authentication process does not check for the next server in the sequence.

"ldap"

When "ldap" is specified, the first Active directory (AD) server is contacted. If the AD server is not reachable, the next AD server is contacted. If the authentication fails, the authentication process does not check for the next server in the sequence.

"tacacs+"

When "tacacs+" is specified, the first Active directory (AD) server is contacted. If the AD server is not reachable, the next AD server is contacted. If the authentication fails, the authentication process does not check for the next server in the sequence.

"radius;local"

Enables the current RADIUS configuration as the primary AAA service and the switch-local database as the secondary AAA service. If "radius" and "local" are specified, and if the RADIUS servers are reachable and the user credentials are correct, the user authentication succeeds. If the user provides credentials from the switch database, the RADIUS authentication fails but login succeeds through the switch database.

"ldap;local"

Enables the current LDAP configuration as the primary AAA service and the switch-local database as the secondary AAA service. If "ldap" and "local" are specified, and if the AD servers are reachable and the user credentials are correct, the user authentication succeeds. If the user provides credentials from the switch database, AD authentication fails but login would still succeed through the switch database.

"tacacs+;local"

Enables the current TACACS+ configuration as the primary AAA service and the switch-local database as the secondary AAA service. If "tacacs+" and "local" are specified, and if the AD servers are reachable and the user credentials are correct, the user authentication succeeds. If the user provides credentials from the switch database, AD authentication fails but login would still succeed through the switch database.

-backup

For use with the "radius;local", "ldap;local", and "tacacs+;local" options only. The backup option states to try the secondary AAA service only if none of the primary AAA services are available.

-nologout

If **-nologout** is not specified, a change in the authentication mechanism may result in the termination of existing sessions. All existing sessions are terminated if the new authentication mechanism is one of the following: RADIUS only, LDAP only, TACACS+ only, or Local. If **-nologout** is specified, there will be no effect on the existing sessions regardless of the chosen authentication mechanism.

--help

Displays the command usage.

Examples

To display the current RADIUS, LDAP, and TACACS+ configurations:

```
switch:admin> aaaconfig --show
RADIUS CONFIGURATIONS
=====
Position Server      Port  Timeout(s) Auth-Protocol
-----
1   192.168.233.48   1812   3         CHAP
2   192.168.233.44   1812   3         CHAP
3   radserver        1812   5         CHAP

LDAP CONFIGURATIONS
=====
LDAP configuration does not exist.

TACACS+ CONFIGURATIONS
=====
TACACS+ configuration does not exist.

Primary AAA Service: Switch database
Secondary AAA Service: None
```

To move the RADIUS server "radserver" from position 3 to position 1:

```
switch:admin> aaaconfig --move radserver -conf radius 1
```

To configure the RADIUS server 192.168.233.48 as an LDAP server:

```
switch:admin> aaaconfig --change 192.168.233.48 \
-conf ldap -p 3002 -s newsecret -t 1
```

To add an AD/LDAP server to the configuration:

```
switch:admin> aaaconfig --add 194.72.68.335 \
-conf ldap -p 3002 -d brocade.com -t 1
```

To replace the AAA service with backup option:

```
switch:admin> aaaconfig --authspec "ldap;local" -backup
```

To change the authentication mechanism with the no logout provision:

```
switch:admin> aaaconfig --authspec radius -nologout
switch:admin> aaaconfig --show
RADIUS CONFIGURATIONS
=====

Position      : 1
Server        : 172.20.1.68
Port          : 1812
Timeout(s)    : 3
Auth-Protocol: PEAP-MSCHAPv2

LDAP CONFIGURATIONS
=====
LDAP configuration does not exist.

TACACS+ CONFIGURATIONS
=====
TACACS+ configuration does not exist.

Primary AAA Service: RADIUS
Secondary AAA Service: None
```

To add a TACACS+ server to the configuration:

```
switch:admin> aaaconfig --add 10.17.56.56 -conf tacacs+ \
-s "sharedkey" -p 49 -t 5 -a pap
switch:admin> aaaconfig --authspec tacacs+
switch:admin> aaaconfig --show
RADIUS CONFIGURATIONS
=====
RADIUS configuration does not exist.

LDAP CONFIGURATIONS
=====
LDAP configuration does not exist.

TACACS+ CONFIGURATIONS
=====
Position      : 1
Server        : 10.17.56.56
Port          : 49
Retries       : 5
```

```
Timeout(s) : 5  
Auth-Protocol: PAP
```

```
Primary AAA Service: TACACS+  
Secondary AAA Service: None
```

SEE ALSO None

ad

ad

Manages Admin Domain operations.

Synopsis

```
ad --activate ad_id
ad --add ad_id [-d "dev_list"] [-s "switch_list"]
ad --apply
ad --clear [-f]
ad --create ad_id [-d "dev_list"] [-s "switch_list"]
ad --deactivate ad_id
ad --delete ad_id
ad --exec ad_id "command_list"
ad --remove ad_id [-d "dev_list"] [-s "switch_list"]
ad --rename ad_id new_ad_id
ad --save
ad --select ad_id
ad --show [-i | [ad_id [-m mode]]] (in AD255 context)
ad --show [-i ] (in AD0 context)
ad --show (in AD1-254 context)
ad --validate [-i | [ad_id | [-m mode]]]
ad --transabort
ad --transshow
```

Description Use this command to manage Admin Domain operations. Admin Domains are not supported in Fabric OS v8.0.1. When this command is executed, it displays the following warning message: "Warning: Admin Domains are not supported in Fabric OS v8.0.1. Admin Domain commands and functionality will be removed in future Fabric OS versions."

This command follows a batched-transaction model. When executed with the **--activate**, **--add**, **--clear**, **--create**, **--deactivate**, **--delete**, **--remove**, **--rename** options, this command changes only the Defined Configuration in the transaction buffer. The **--save** option sends the changes made in the transaction buffer to all other switches and permanently saves the changes to the Defined configuration in persistent storage. The **--apply** option performs a save operation, sends a request to apply the Admin Domain configuration (as defined in the persistent storage), and then enforces the configuration locally.

The Admin Domain transaction buffer is linked to the current login shell and is lost on logout. Use the **--transshow** option to display the current Admin Domain transaction information.

Before creating Admin Domains, the default zone mode should be set to "No Access". To set the default zone mode to "No Access" execute the following command sequence:

```
switch:admin> ad --select ADO
switch:admin> defzone --noaccess
switch:admin> cfgsave
```

Refer to **defZone** help for more information.

All switches, switch ports and devices in the fabric that are not specified in any other Admin Domain are treated as implicit members of ADO. Members added to ADO are called explicit members.

When a new Admin Domain is created, the members included in the new Admin Domain are automatically removed from the implicit member list of ADO. If the devices included in the new Admin Domain are already zoned in ADO, and if you want to move these devices from ADO without any traffic disruption, do the following:

1. Add the devices to ADO's explicit member list using **ad --add** and **ad --apply**.
2. Create new ADs with the devices and execute **ad --apply**.
3. Select (or login to) the new Admin Domain and create a relevant zone configuration and zones (Refer to **zone --copy** help for details). Enable the new zone configuration under the Admin Domain.
4. (Optionally) remove explicit members from ADO (using **ad --remove** and **ad --apply**). Remove the member references from the ADO zone database.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

The generalized syntax for this command is "**ad action arguments**". Command *actions* are described first. Argument details follow the description of actions.

The following *actions* are supported:

--activate arg

Activates an Admin Domain. An Admin Domain is enforced only if it is in an activated state. ADO is always in an activated state. By default, after the Admin Domain is enabled, the devices specified in the Admin Domain are not able to see each other until they are zoned together. The command prompts for confirmation. The activate operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

--add arguments

Adds new members to an existing Admin Domain. The add operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

--apply

Saves the current transaction buffer contents to the defined configuration and enforces the defined configuration on all switches in the fabric. **ad --apply** prompts for confirmation.

--clear [-f]

Deletes all ADs and clears ADO's explicit members. This command fails if AD1 through AD254's zone databases are not empty. The command prompts for confirmation before deleting all Admin Domains. The clear operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

When used with the **-f** option, this command deletes all ADs and clears all explicit members from ADO, even if the zone databases or AD1 through AD254 are not empty. Use the force option to remove Admin Domains before enabling Virtual Fabrics.

--create arguments

Creates a new Admin Domain with optionally specified device or switch members. A newly created Admin Domain is in an activated state. It initially contains no zone database. The newly created Admin Domain remains in the transaction buffer until you issue **ad --apply** or **ad --save**. ADO always exists. Use **ad --add** to add explicit members to ADO.

--deactivate arg

Deactivates an Admin Domain. This operation fails if an effective zone configuration exists under the Admin Domain. This operation is not allowed on ADO. **ad --deactivate** does not disable any ports. Existing logins to a deactivated Admin Domain are not terminated; however, subsequent CLI execution is disallowed. A message is displayed to indicate that the current Admin Domain is not active. The command prompts for confirmation. The deactivate operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

--delete arg

Deletes an Admin Domain. This command succeeds regardless of whether the Admin Domain is in a deactivated or an activated state. ADO always exists; using this operation on ADO does not delete ADO, it only removes all explicit members from ADO. The ADO zone database does not need to be empty for the delete operation to succeed. Not all existing user sessions to a deleted Admin Domain are terminated; however, subsequent CLI execution is disallowed. A message displays indicating that the current Admin Domain is not active. The command prompts for confirmation before executing the delete action. The delete operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

--exec arguments

This command performs the following tasks:

- Creates a new shell.
- Executes **ad --select** to the specified Admin Domain.
- Executes the specified commands.
- Exits the shell.

--remove arguments

Removes one or more members from an Admin Domain. Removing the last member from an Admin Domain deletes the Admin Domain. The remove operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

--rename arguments

Renames the specified Admin Domain. If a reserved name is used for *new_ad_id* (AD number format), the operation fails if the reserved name does not correspond to the *ad_id* AD number. The rename operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

--save

Saves the outstanding Admin Domain transaction to the defined configuration on all switches in the fabric. The saved Admin Domain definition is enforced only when **ad --apply** is issued. Attempts to modify and save an Admin Domain that is currently enforced will fail. The command prompts for confirmation.

--select *arg*

Selects an Admin Domain context. This command fails if the corresponding Admin Domain is not activated. This operation succeeds only if you have the specified Admin Domain. This command internally spawns off a new shell within the requested Admin Domain context. Type **logout** or **exit** to exit from the selected Admin Domain. The zone transaction is linked to the current shell; therefore, the zone transaction buffer is lost on logout. Use **cfgTransShow** to display the current zoning transaction information.

--show *arguments*

Displays the membership information of the specified Admin Domain or all Admin Domains.

When executed in an AD255 context and an Admin Domain name is not specified, all information about all existing Admin Domains is displayed. When executed in an ADO-AD254 context, the command, by default, displays the members of the current Admin Domain's effective configuration, and therefore you cannot specify an *ad_id* or *mode*.

When executed in an AD255 context, all Admin Domain information from the transaction buffer, defined configuration and effective configuration is displayed.

--validate *arguments*

Checks whether Admin Domain members are from a non-Admin Domain aware switch or the members do not exist in the fabric. The output is similar to **ad --show**; however, all members that are from non-Admin Domain aware switches are marked with a plus sign (+). Members that are not online are marked with an asterisk (*).

FC Router Front Phantom Domain and FC Router Translate Phantom Domain are virtual entities without any exposed management interfaces; therefore, any FC Router phantom switch WWN specified in an AD switch member list is marked as a non-Admin Domain aware member. All **D,PI** members in the device list corresponding to an FC Router Phantom Domain are marked as non-Admin Domain aware members. All FC Router imported devices in the AD device list are marked as AD-aware members.

--transabort

Aborts the transaction buffer. The command prompts for confirmation before aborting the transaction.

--transshow

Displays the ID of the current Admin Domain transaction and indicates whether or not the transaction can be aborted. The transaction cannot be aborted if it is an internal Admin Domain transaction.

The following *arguments* are supported with selected AD *actions*:

ad_id

Uniquely identifies an Admin Domain. An *ad_id* can be a *name* or a *number*.

| | |
|-------------------------|--|
| <i>name</i> | An Admin Domain name can be up to 63 bytes, must begin with a letter, and can consist of letters, numbers, and underscore characters. The Admin Domain names with the format AD[0-255] are reserved for autoassigning Admin Domain names to Admin Domains created with an Admin Domain number, and can be assigned only to the corresponding Admin Domain. Using ad --rename , for example, in an attempt to assign a name of AD5 to an Admin Domain with ID not equal to 5 fails. Admin Domain names are case-sensitive. |
| <i>number</i> | An Admin Domain can be specified by a number. Valid values include 0 through 255. ADO and AD255 are always active. ADO cannot be specified with --activate , --deactivate or --delete actions. AD255 can be specified only with --exec , --show and --validate actions. For all command iterations, with the exception of ad --create , the Admin Domain is specified either by a name or a number. For ad --create , both name and number can be specified: for example, ad --create test_ad/10 -d "100,5; 100,1" . |
| -d "dev_list" | Specifies the list of devices in an Admin Domain, in quotation marks. Separate each entry in the device list with a semicolon (;). Valid formats include the following: |
| <i>D,PI</i> | Uses existing zone <i>D,PI</i> member types. Benefits include the following: <ul style="list-style-type: none"> • Grants port control and zoning on the switch port and the devices attached to that port. • <i>PI</i> can be specified as a range; for example, <i>D,[0 to 34]</i>. The port index range is expanded and stored internally. • The same <i>D,PI</i> members can be specified in more than one Admin Domain. |
| <i>Device WWN</i> | Uses traditional zone WWN member types. Benefits include the following: <ul style="list-style-type: none"> • Supports node or port WWNs. • End-device members, whose WWNs are used in an Admin Domain definition, need not be online when the Admin Domain is created (similar to a zoneCreate operation). • Provides rights to zone the devices. • Provides administrative view rights to the switch port the device is connected to. |
| -s "switch_list" | Specifies the list of switches in an Admin Domain. The list must be enclosed in quotation marks. Separate each entry in the switch list with a semicolon (;). Specify the switch in one of the following formats: |
| <i>Switch WWN</i> | World wide name of the switch. |
| <i>Domain ID</i> | Any switch member specified in Domain ID format is converted into a switch WWN-based on the current fabric information. Operations with switch list fail if the domain ID to switch WWN lookup fails. Membership in an AD <i>switch_list</i> grants switch administrative operations such as switchDisable , switchEnable , reboot , ad , etc. on the switch. Ownership of a switch implicitly provides port control capability on all its ports, but no zoning control. |

| | |
|-----------------------------|--|
| <code>"command_list"</code> | Specifies one or more commands to execute in an Admin Domain context. This operand is valid only with the <code>--exec</code> option. |
| <code>new_ad_id</code> | Specifies a new Admin Domain name or number. This operand is valid only with the <code>--rename</code> option. Format is the same as <code>ad_id</code> . |
| <code>-i</code> | Displays the implicit members of ADO. This operand is valid only with the <code>--show</code> option. |
| <code>-m mode</code> | Specifies the mode in which Admin Domain configuration information is displayed. This operand is valid only with <code>--show</code> and <code>--validate</code> . Valid values for <code>mode</code> include the following: |
| <code>0</code> | Displays the Admin Domain configuration in the current transaction buffer. |
| <code>1</code> | Displays the Admin Domain configuration stored in persistent memory (defined configuration). |
| <code>2</code> | Displays the currently enforced Admin Domain configuration currently enforced (effective configuration). |

Examples

To enable AD5:

```
switch:admin> ad --activate 5
You are about to activate a new admin domain.
Do you want to activate '5' admin domain \
  (yes, y, no, n): [no] y
```

To enable AD_13:

```
switch:admin> ad --activate AD_13
```

To add new device members to AD1:

```
switch:admin> ad --add AD1, -d "100,5; 4,1"
```

To apply all changes made to the Admin Domain configurations since `--apply` was last executed:

```
switch:admin> ad --apply
You are about to enforce the saved AD configuration.
This action will trigger ad --apply to all switches \
in the fabric.
Do you want to apply all admin domains \
  (yes, y, no, n): [no] y
```

To clear all Admin Domain definitions:

```
switch:admin> ad --clear
You are about to delete all ADs definitions.
This operation will fail if zone configurations \
exists in AD1-AD254
Do you want to clear all admin domains \
(yes, y, no, n): [no] y
```

To create an Admin Domain with a mix of D,PI, WWNs, and zone alias device members (two different methods shown):

```
switch admin> ad --create "AD1", \
-d "100,5; 1,3; 20:00:00:e0:8b:05:4d:05"

switch admin> ad --create 1, \
-d "100,5; 1,3; 21:00:00:e0:8b:05:4d:05"
```

To create an Admin Domain with two switches identified by domain ID and switch WWN:

```
switch:admin> ad --create "AD1", \
-s "100; 10:00:00:60:69:80:59:13"
```

To create an Admin Domain with a device list and a switch list:

```
switch:admin> ad --create "AD1", \
-d "100,5; 1,3; 21:20:00:00:e0:8b:05:4d:05" \
-s "100; 10:00:00:60:69:80:59:13"
```

To deactivate Admin Domain 5:

```
switch:admin> ad --deactivate 5
You are about to deactivate an AD.
This operation will fail if an effective zone \
configuration exists in the AD
Do you want to deactivate '5' admin domain \
(yes, y, no, n): [no] y
```

To delete AD13:

```
switch:admin> ad --delete 13
You are about to delete an AD.
This operation will fail if an effective zone \
configuration exists in the AD
Do you want to delete '13' admin domain \
(yes, y, no, n): [no] y
```

To execute **switchShow** in an AD7 context (using the current *user_id*):

```
switch:admin> ad --exec 7 "switchshow"
```

To rename Eng_AD to Eng_AD2:

```
switch:admin> ad --rename Eng_AD Eng_AD2
```

To rename AD 200 to Eng_AD200:

```
switch:admin> ad --rename 200 Eng_AD200
```

To rename a user-assigned Admin Domain name to a reserved Admin Domain name (this operation fails if AD_test's AD number is not 200):

```
switch:admin> ad--rename AD_test AD200
```

To remove the devices 100,5 and 1,3 from AD1:

```
switch:admin> ad --remove "AD1", \
-d "100,5; 1,3; 21:00:00:e0:8b:05:4d:05"
```

To remove the switch 100 from AD1:

```
switch:admin> ad --remove "AD1", -s "100"
```

To save any outstanding Admin Domain definition-related transaction buffer

```
switch:admin> ad --save
You are about to save the outstanding AD membership.
This action will only save the changes to Defined \
configuration.
Any changes made will be enforced only on ad --apply.
Do you want to save admin domains \
(yes, y, no, n): [no] y
```

To select a new Admin Domain context by specifying the AD number:

```
switch:admin> ad --select 12
```

To display all ADs:

```
switch:admin> ad --show
Current AD: 255 : AD255

Transaction buffer configuration:
-----
no configuration

Defined configuration:
-----

AD: 1 :          AD1      Active

Device WWN members:  21:00:00:80:e5:12:8b:37;
                    21:00:00:80:e5:12:8b:55;
Switch port members: 1,0; 1,1; 1,2; 1,3; 1,4; 1,5;
                    1,6; 1,7; 1,8; 1,9; 1,10; 1,11;
                    1,12; 1,13; 1,14; 1,15;
Switch WWN members:  10:00:00:60:69:00:02:53;

Effective configuration:
-----

AD: 1 :          AD1      Active

Device WWN members:  21:00:00:80:e5:12:8b:37;
```

ad

```
                21:00:00:80:e5:12:8b:55;
Switch port members: 1,0; 1,1; 1,2; 1,3; 1,4; 1,5;
                    1,6; 1,7; 1,8; 1,9; 1,10; 1,11;
                    1,12; 1,13; 1,14; 1,15;
Switch WWN members: 10:00:00:60:69:00:02:53;
```

To display the AD1 configuration information in the transaction buffer:

```
switch:admin> ad --show 1 -m 0
Current AD: 255 : AD255

Transaction buffer configuration:
-----
no configuration
```

To display the AD10 configuration information in persistent storage:

```
switch:admin> ad --show 10 -m 1
Current AD: 255 : AD255

Defined configuration:
-----

AD: 1 :          AD1      Active

Device WWN members: 21:00:00:80:e5:12:8b:37;
                   21:00:00:80:e5:12:8b:55;
Switch port members: 1,0; 1,1; 1,2; 1,3; 1,4; 1,5;
                   1,6; 1,7; 1,8; 1,9; 1,10; 1,11;
                   1,12; 1,13; 1,14; 1,15;
Switch WWN members: 10:00:00:60:69:00:02:53;
```

To display the Admin Domain effective configuration information:

```
switch:admin> ad --show -m 2
Current AD: 255 : AD255

Effective configuration:
-----

AD: 1 :          AD1      Active

Device WWN members: 21:00:00:80:e5:12:8b:37;
                   21:00:00:80:e5:12:8b:55;
Switch port members: 1,0; 1,1; 1,2; 1,3; 1,4; 1,5;
                   1,6; 1,7; 1,8; 1,9; 1,10; 1,11;
                   1,12; 1,13; 1,14; 1,15;
Switch WWN members: 10:00:00:60:69:00:02:53;
```

To display the configuration information in the transaction buffer:

```
switch:admin> ad --validate
Current AD Number: 255 AD Name: AD255

Transaction buffer configuration:
-----
no configuration
```

Defined configuration:

AD Number: 1 AD Name: AD1 State: Inactive

```

Device WWN members:      10:00:00:00:00:01:00:00;
                          10:00:00:00:00:00:04:00:00;
                          10:00:00:00:00:00:05:00:00;
                          10:00:00:00:00:00:06:00:00;
                          10:00:00:00:00:00:08:00:00;
                          10:00:00:00:00:00:03:00:00;
                          10:00:00:00:00:00:02:00:00;
                          10:00:00:00:00:00:07:00:00;
                          10:00:00:00:00:00:15:00:00;
                          10:00:00:00:00:00:16:00:00;
                          10:00:00:00:00:00:17:00:00;
                          10:00:00:00:00:00:18:00:00;
                          10:00:00:00:00:00:11:00:00;
                          10:00:00:00:00:00:12:00:00;
                          10:00:00:00:00:00:13:00:00;
                          10:00:00:00:00:00:14:00:00;

```

Effective configuration:

AD Number: 50 AD Name: AD50 State: Active

```

Device WWN members:      10:00:00:00:00:00:17:00:00;
                          10:00:00:00:00:00:15:00:00;
Switch port members:     2,52; 2,53; 2,54; 2,55; 21,5;
                          3,28; 3,29; 98,72; 98,75;
                          69,16; 69,18; 69,21; 1,336;
                          1,337;

```

```

-----
* - Member does not exist
+ - Member is AD Unaware

```

To abort the Admin Domain management transaction buffer:

```

switch:admin> ad --transabort
You are about to abort the outstanding AD transaction.
Do you want to abort the AD transaction \
(yes, y, no, n): [no] y

```

To display the current Admin Domain transaction:

```

switch:admin> ad --transshow
Current transaction token is 26816
It is abortable

switch:admin> ad --transshow
There is no outstanding zoning transaction

```

SEE ALSO [cfgSave](#), [cfgTransShow](#), [defZone](#), [logout](#)

ag

ag

Enables Access Gateway (AG) and manages AG-specific operations.

Synopsis

```
ag --help
ag --show
ag --modeshow | --modeenable | --modedisable
ag [--policyenable | --policydisable] policy
ag --policyshow
ag --mapshow [N_Port]
ag [--mapset | --mapadd | --mapdel] N_Port [F_Port1; F_Port2;...]
ag [--staticadd | --staticdel] N_Port [F_Port1; F_Port2;...]
ag --pgshow [pgid]
ag --pgcreate pgid "N_Port1 [N_Port2;...]" [-n pgname]
    [-m "lb; mfnm"]
ag [--pgadd | --pgdel] pgid "N_Port1 [N_Port2;...]"
ag --pgrename pgid newname
ag --pgremove pgid
ag [--pgmapadd | --pgmapdel] pgid "F_Port1 [F_Port2;...]"
ag [--pgsetmodes | --pgdelmodes] pgid "mfnm;lb"
ag --pgfnmtov [new_tov]
ag [--failoverenable | --failoverdisable] [N_Port | -pg pgid]
ag --failovershow [N_Port]
ag [--failbackenable | --failbackdisable] [N_Port | -pg pgid]
ag --failbackshow [N_Port]
ag [--prefset | --prefdel] "F_Port [F_Port2;...]" N_Port
ag --prefshow
ag [--adsset | --adsadd | --adsdel] "F_Port [F_Port2;...]"
    "WWN [WWN2;...]"
ag --adsshow
ag --persistentalpaenable 1 | 0 mode
```

```
ag --printalpamap F_Port
ag --deletepwwnfromdb PWWN
ag --clearalpamap F_Port
ag --addwwnmapping N_Port "WWN [WWN2;...]" | --all
ag --delwwnmapping N_Port "WWN [WWN2;...]" | --all
ag --addwwnpgmapping Port_Group "WWN [WWN2;...]" | --all
ag --delwwnpgmapping Port_Group "WWN [WWN2;...]" | --all
ag --addwwnfailovermapping N_Port "WWN [WWN2;...]" | --all
ag --delwwnfailovermapping N_Port "WWN [WWN2;...]" | --all
ag --wwnmappingenable "WWN [WWN2;...]" | --all
ag --wwnmappingdisable "WWN [WWN2;...]" | --all
ag --wwnmapshow
ag --reliabilitycounterset
ag --reliabilitycountershow
ag --backupmappingsave N_Port
ag --backupmappingdel N_Port
ag --backupmappingshow N_Port
```

Description Use this command to perform the following Access Gateway management functions:

- Enable or disable Access Gateway mode.
- Display current configuration and state of AG.
- Configure and display F_Port to N_Port mapping.
- Configure and display N_Port failover and failback policies.
- Configure and display Port Group policy.
- Create or remove a Port group.
- Enable or disable auto port configuration (APC) modes for a Port group.
- Get or set timeout value for fabric name monitoring.
- Display Port Groups and member N_Ports.
- Add or delete N_Ports in a Port group.
- Display all policies and their status.
- Enable or disable auto port configuration (APC) policy.
- Enable or disable preferred secondary N_Port policy.
- Enable, disable, and manage advanced device security (ADS) policy.
- Manage persistent ALPA mode.
- Manage device WWN to N_Port mappings.
- Manage device WWN to N_Port group mappings.
- Manage device WWN failover to N_Ports configured as preferred failover ports.
- Enable or disable device WWN mappings.
- Configure the reliability limit for the preferred N_Port.
- Save the configured F_Ports, static F_Ports for the given N_Port, and F_Ports for which the given N_Port is a preferred one.
- Delete the backup mappings for the given N_Port, if any.
- Display the saved mappings for the given N_Port, if any.

AG configuration changes are saved persistently as configuration keys. Use the **portCfgNPort** command to set a port as N_Port.

This command supports multiple configurations for mapping device logins to N_Ports for the purposes of load balancing and redistribution in the event of a fabric change. If multiple mappings are configured, the system considers the available mappings in a fixed order of priority to determine which of the available N_Ports should be assigned to the login request. The first eligible mapping is chosen in the order specified below.

1. Device WWN to N_Port
2. Device WWN to N_Port Group
3. Automatic device WWN load balancing.
4. F_Port to N_Port
5. F_Port to N_Port Group

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

AG is supported only on selected Brocade hardware platforms. Refer to the *Access Gateway Administrator's Guide* for hardware support and AG configuration procedures.

In non-AG mode, only two actions are available: **--modeenable** and **--modeshow**.

The **--backupmappingsave**, **--backupmappingdel** and **--backupmappingshow** commands can be executed if the auto policy is disabled.

Operands This command has the following operands:

--help

Displays the command usage.

--show

Displays the current configuration of the Access Gateway. This includes all N_Ports and F_Ports that are currently online, failover and failback settings as well as any online F_Ports that are currently mapped to N_Ports. Failover and failback policies are displayed as enabled (1) or disabled (0).

--modeshow

Displays the current Access Gateway operating mode of the switch as either enabled or disabled.

--modeenable

Enables Access Gateway mode on a switch. Long distance mode settings should be cleared for all ports on the NPIV edge switch to which the AG is connected. Otherwise, the NPIV switch port displays the long distance port type along with the F_Port.

--modedisable

Disables Access Gateway mode on a switch. After AG mode is disabled, the switch reboots automatically and comes online with default zone access set to "No Access". In order to merge the switch to a fabric, set the default zone to "All Access" and disable/enable the E_Port.

--policyshow

Displays the supported AG port policies and their status as either enabled or disabled. AG supports four types of policies:

- Port Grouping (pg) policy: This policy manages failover of an F_Port to a set of related N_Ports in a port group.
- Auto Port Configuration (auto): When this policy is enabled, the AG enabled switch automatically detects available ports and maps F_Ports to N_Ports. Auto port configuration is disabled by default.
- Advanced Device Security (ADS) policy. This policy restricts access to the fabric at the AG level to a set of authorized devices. Unauthorized access is rejected and the system logs a RASLOG message. You can configure the list of allowed devices for each F_Port by specifying their Port WWN. Refer to the **ag --ads*** commands for information on managing advanced device security. ADS policy is disabled by default, which means that all devices can connect to the switch.
- WWN Based Load Balancing policy: This policy routes device logins to the least loaded port in the port group to which they are mapped.

--policyenable *policy*

Enables the specified port policy for the Access Gateway. When a new policy is enabled, all port related configuration settings are lost. Use the **configUpload** command to save the current port configuration. Valid policies include the following:

pg

Enables the port grouping policy. A default port group "pg0" is created, which includes all configured N_Ports assigned to the policy. Enabling port grouping policy disables the Get Fabric Name policy.

auto

Enables the automatic port configuration policy. When enabled, this policy applies to all ports on the switch. All F_Port to N_Port mapping and port group configurations are ignored.

ads

Enables the advanced device security (ADS) policy. When enabled, this policy applies to all the ports on the switch. By default all devices have access to the fabric on all ports.

wwnloadbalance

Enables the device WWN load balancing policy. When this policy is enabled, device logins are sent to the least loaded port in the port group to which they are mapped. These devices are displayed with **ag --wwnmapshow** as dynamic device mappings. Port Group policy must be enabled before you can enable the WWN load balancing policy.

--policydisable *policy*

Disables the specified policy for the Access Gateway. When a policy is disabled, all port-related configuration settings are lost. Use the **configUpload** command to save the current port configuration. Valid policies include the following:

pg

Disables the port grouping policy. All port group configurations are deleted. Disabling port grouping policy enables the Get Fabric Name policy.

auto

Disables the automatic port configuration policy and deletes all associated configuration settings.

ads

Disables the advanced device security (ADS) policy and deletes all lists of allowed device WWNs.

wwnloadbalance

Disables the device WWN load balancing policy.

--mapshow [*N_Port* | *device_WWN*]

Displays the F_Ports that are configured and currently mapped to a given "primary" N_Port. Optionally specify an N_Port to display the F_Ports that are mapped to the specified N_Port only, or specify a device WWN to display the N_Port to which the device WWN is mapped. Failover and failback policies are displayed as enabled (1) or disabled (0).

--mapset *N_Port* [*F_Port1*; *F_Port2*;...]

Maps a set of F_Ports to a specified "primary" N_Port forcing all traffic from the F_Ports to be routed through this N_Port to the attached fabric. An F_Port cannot be mapped to more than one primary N_Port at any given time. F_Ports are enabled only if the N_Port is online. This command overwrites existing port mappings. Use a blank list ("") to clear current mappings.

--mapadd *N_Port* *F_Port1* [; *F_Port2*;...]

Adds one or more specified F_Ports to the mapping of an existing "primary" N_Port. The traffic for the configured F_Ports is routed to the fabric through the specified N_Port when the F_Ports come online. An F_Port cannot be mapped to more than one primary N_Port at the same time.

--mapdel *N_Port* *F_Port1* [; *F_Port2*;...]

Deletes one or more specified F_Ports from the "primary" N_Port mapping.

--staticadd *N_Port* *F_Port1* [; *F_Port2*;...]

Creates a static mapping between an existing "primary" N_Port and one or more specified F_Ports. This command removes existing mappings. Once the static mapping is enabled, the F_Ports and all logged-in devices will logout of the previous N_Port and login using the new N_Port.

--staticdel *N_Port* *F_Port1* [; *F_Port2*;...]

Deletes a static mapping between an existing "primary" N_Port and one or more specified F_Ports. This command removes existing mappings. Alternately, you can remove an existing mapping by mapping the F_Port to another N_Port.

--pgshow [*pgid*]

Displays the Port Group configuration. The port grouping feature supports specifying a set of N_Ports to be included in the Port Group (PG) Policy. The factory default PG is "pg0", which includes all N_Ports. The default PG cannot be removed or renamed.

--pgcreate *pgid*

"N_Port1 [;*N_Port2*;...]" [-n *pgname*][-m "lb; mfnm"]

Creates a port group with the ID *pgid* and a specified list of N_Ports to be included in the policy. The list must be enclosed in quotation marks. Ports must be separated by semicolons. The Port Group ID must not exceed 64 characters. Optionally specify a name for the port group and a mode. Modes are by default disabled. For an explanation of mode values, refer to **--pgsetmodes**.

--pgadd *pgid* "*N_Port1* [; *N_Port2*;...]"

Adds one or more N_Ports to the specified port group. The port list must be enclosed in quotation marks. Ports must be separated by semicolons.

- pgdel *pgid* "*N_Port1* [*; N_Port2*;...]"**
 Deletes one or more *N_Ports* from the specified port group. Deleted ports are added to the default port group "pg0". The port list must be enclosed in quotation marks. Ports must be separated by semicolons.
- pgrename *pgid newname***
 Replaces the name of an existing port group with the specified new name. The port group ID must not exceed 64 characters.
- pgremove *pgid***
 Deletes the specified port group. The *N_Ports* in the port group that was deleted are moved to the default port group, which is *pgid 0*.
- pgmapadd *pgid* "*F_Port1* [*;F_Port2*;...]"**
 Maps the specified *F_Ports* to the PG identified by the *pgid*. Upon execution, the system identifies the least loaded *N_Port* in the port group and maps the *F_Ports* to that *N_Port*. The port list must be enclosed in double quotation marks. Ports must be separated by semicolons. Login balancing (LB) mode must be enabled on the port group for this command to succeed. Use **ag --pgsetmodes** to enable LB mode.
- pgmapdel *pgid* "*F_Port1* [*;F_Port2*;...]"**
 Removes one or more *F_Ports* that are part of the port group identified by the *pgid* from their mapping to a corresponding *N_Port*. The port list must be enclosed in double quotation marks. Ports must be separated by semicolons. Login balancing (LB) mode must be enabled on the port group for this command to succeed. Use **ag --pgsetmodes** to enable LB mode.
- pgsetmodes *pgid* "*lb*;*mfnm*"**
 Sets the APC modes for the specified port group. The mode list must be enclosed in double quotation marks and the modes must be separated by a semicolon. Alternately you can set the modes at the time when you create the port group with the **pgcreate** command. The following modes are supported:
- lb**
 Specifies the login balancing mode for the specified port group. If login balancing mode is enabled and an *F_Port* goes offline, logins in the port group are redistributed among the remaining *F_Ports*. Similarly, if an *N_Port* comes online, port logins in the PG are redistributed to maintain a balanced *N_Port* to *F_Port* ratio. This operation is disruptive. Login balancing mode is disabled by default in all port groups.
- mfnm**
 Enables the managed fabric name monitoring mode (MFNM) in the specified port group. This command changes the fabric name monitoring mode from "default" to "managed". In both default and managed mode, the system queries the fabric name once every 120 seconds, and if it detects an inconsistency, for example, if the port group is connected to multiple fabrics, it triggers a RASLOG message. The difference between default and managed fabric name monitoring is that in managed mode, failover is disabled for all ports in the port group if the system detects an inconsistency in fabric names.
- You can add or remove MFNM from a port group, however, doing so will enable or disable MFNM on the entire switch. RASLOG messages are generated only if MFNM is enabled on the entire switch and multiple fabrics are connected to the switch.

--pgdelmodes *pgid* "lb;mfnm"

Disables the specified modes on a given port group. The mode list must be enclosed in double quotation marks and the modes must be separated by a semicolon. For a description of supported modes, refer to **--pgsetmodes**.

--pgfnmtov *new_tov*

Displays the fabric name monitoring timeout value in seconds when used without specifying a new value. To change the current value, specify a new timeout value in seconds. The valid range is 30 to 120 seconds. The default value is 120 seconds.

--failoverenable [*N_Port*] -pg *pgid*

Enables the failover policy for a given N_Port or for all N_Ports in the given port group. When failover policy is enabled for a given N_Port, F_Ports behave as follows:

- If only primary F_Port to N_Port mapping is in place, all currently mapped F_Ports fail over to another available N_Port in the event the original N_Port becomes disabled. If multiple N_Ports are available for failover, F_Ports are evenly balanced across all available N_Ports. If no other N_Port is available, failover does not occur.
- If preferred secondary F_Port to N_Port Mapping is in place, the F_Ports are routed through the preferred Secondary N_Port. If the preferred secondary N_Port is offline, the F_Ports are disabled.

--failoverdisable [*N_Port*] -pg *pgid*

Disables the failover policy for a given N_Port or for all N_Ports in the given port group.

--failovershow [*N_Port*]

If an N_Port is specified (optional), the command displays the failover policy for this N_Port. Otherwise, the failover policy for all N_Ports is displayed. Failover is displayed as enabled (1) or disabled (0).

--failbackenable [*N_Port*] -pg *pgid*

Enables the failback policy for a specified N_Port or for all N_Ports in the given port group. When failback policy is enabled, ports behave as follows:

- If only primary F_Port to N_Port mapping is in place, all F_Ports are automatically rerouted back to the N_Ports to which they were originally mapped as those N_Ports come back online. Only the originally mapped F_Ports fail back. In the case of multiple N_Port failures, only F_Ports that were mapped to the recovered N_Port experience failback. The remaining F_Ports are not redistributed among the online N_Ports during the failback.
- If preferred secondary F_Port to N_Port mapping is in place, and the primary N_Port comes back online, then the F_Ports are rerouted through the primary N_Port. If the secondary N_Port comes online, while the primary N_Port is still offline, F_Ports are rerouted through the secondary N_Port.

--failbackdisable [*N_Port*] -pg *pgid*

Disables the failback policy for the specified N_Port or for all N_Ports in the given port group.

--failbackshow [*N_Port*]

If an N_Port is specified (optional), the command displays the failback policy for this N_Port. Otherwise, the failover policy for all the N_Ports is displayed. The failback policy is displayed as disabled (0) or enabled (1).

- prefset "F_Port[:F_Port2;...]" N_Port**
Sets the preferred secondary N_Port for one or more F_Ports. Preferred mapping is optional. Preferred F_Port to N_Port Mapping provides an alternate N_Port for F_Ports to come online for predictable failover and failback. An F_Port must have primary N_Port mapping before a secondary N_Port can be configured. The list of F_Ports to be mapped must be enclosed in double quotation marks. Port numbers must be separated by semicolons.
- prefdel "F_Port[:F_Port2;...]" N_Port**
Deletes the preferred Secondary N_Port for the specified F_Ports. The list of F_Ports to be deleted from the secondary mapping must be enclosed in double quotation marks. Port numbers must be separated by semicolons.
- prefshow**
Displays the preferred Secondary N_Port for all F_Ports.
- adsset "F_Port[:F_Port2;...]" "WWN[:WWN2;...]"**
Sets the list of devices that are allowed to login to a specified set of F_Ports. Devices are specified by their world wide names. Lists must be enclosed in double quotation marks. List members must be separated by semicolons. The maximum number of entries in the allowed device list is twice the per port maximum login count. Replace the WWN list with an asterisk (*) to indicate all access on the specified F_Port list. Replace the F_Port list with an asterisk (*) to add the specified WWNs to all the F_Ports' allow lists. A blank WWN list ("") indicates no access. ADS policy must be enabled for this command to succeed.
- adsadd "F_Port[:F_Port2;...]" "WWN[:WWN2;...]"**
Adds the specified WWNs to the list of devices allowed to login to the specified F_Ports. Lists must be enclosed in double quotation marks. List members must be separated by semicolons. Replace the F_Port list with an asterisk (*) to add the specified WWNs to all the F_Ports' allow lists. ADS policy must be enabled for this command to succeed.
- adsdel "F_Port[:F_Port2;...]" "WWN[:WWN2;...]"**
Deletes the specified WWNs from the list of devices allowed to login to the specified F_Ports. Lists must be enclosed in double quotation marks. List members must be separated by semicolons. Replace the F_Port list with an asterisk (*) to remove the specified WWNs from all the F_Ports' allow lists. ADS policy must be enabled for this command to succeed.
- adsshow**
Displays the list of allowed device WWNs for all F_Ports.
- persistentalpaenable 1 | 0 mode**
Configures the persistent ALPA feature. Once enabled, the ALPA parts of all device PIDs become persistent regardless of whether they were logged in before or after the persistent ALPA feature was enabled. ALPA persistence ensures that there is no inconsistency between logged in devices. The persistent ALPA feature is disabled by default.
- 1 | 0**
Specify 1 to enable persistent ALPA. Specify 0 to disable the feature.

- mode*
- Specifies the manner in which the ALPA is obtained in the event that the ALPA value is already taken by another host. Valid modes include the following:
- s** Specifies a stringent ALPA request mode. In stringent mode, the login is rejected if the ALPA is not available.
 - f** Specifies a flexible ALPA request mode. In flexible mode, the host login is accepted either with the requested ALPA value or with a different ALPA value if the requested ALPA is not available.
- printalpamap *F_Port***
Displays the database entry for the specified port. An *F_Port* must be specified. The output displays the PWWN-to-host-ALPA mapping.
- deletepwwnfrommdb *PWWN***
Removes the specified port WWN entry from the database after the host has logged out.
- clearalpamap *F_Port***
Clears the ALPA values for the specific *F_Port*. This command removes the PWWN-to-ALPA-value mapping from the database.
- addwwnmapping *N_Port* "*WWN* [*WWN2*;...]" | **--all****
Maps one or more device WWNs to a preferred *N_Port*. All traffic from the specified devices is forced through the specified *N_Port*, regardless of which *F_Port* the device logs into. In the event the designated *N_Port* should become unavailable, an alternate port can serve as a preferred failover port. This command only affects devices that are connecting to the fabric after successful execution of this command; it will not affect devices already logged in. If a device is already connected to the switch when its mapping is created, that mapping goes into effect the next time the device connects. The WWN list must be enclosed in double quotation marks. WWNs must be separated by semicolons.
- The **--all** option indicates all device WWNs already mapped, for example, if you wish to change an existing WWN mapping. It does not affect device WWNs that are not part of an existing mapping.
- delwwnmapping *N_Port* "*WWN* [*WWN2*;...]" | **--all****
Removes the mapping of one or more device WWNs to a preferred *N_Port*. The **--all** option removes the mapping for all device WWNs currently mapped to the specified *N_Port*. The WWN list must be enclosed in double quotation marks. WWNs must be separated by semicolons. The mappings are removed upon execution of this command.
- addwwnpgmapping *PG* "*WWN* [*WWN2*;...]" | **--all****
Maps one or more device WWNs to any of the *N_Ports* included in the specified port group . The port group is identified by its port group ID. The **--all** option maps all currently mapped device WWNs to the specified port group. The WWN list must be enclosed in double quotation marks. WWNs must be separated by semicolons. The device WWN to port group mapping takes effect the next time the device logs in.

- delwwnpgmapping** *PG* "*WWN* [*WWN2*;...]" | **--all**
 Removes the mapping between the specified device WWNs and the specified port group. The port group is identified by its port group ID. The **--all** option removes the mapping of all device WWNs currently mapped to the specified port group. The WWN list must be enclosed in double quotation marks. WWNs must be separated by semicolons. The mappings are removed upon execution of this command.
- addwwnfailovermapping** *N_Port* "*WWN* [*WWN2*;...]" | **--all**
 Maps one or more device WWNs to a preferred failover *N_Port*. If the *N_Port* to which the WWNs are mapped is not available or goes down, the device logins fail over to the preferred failover *N_Port*. The **--all** option maps all currently mapped device WWNs to the specified failover *N_Port*. The WWN list must be enclosed in double quotation marks. WWNs must be separated by semicolons.
- delwwnfailovermapping** *N_Port* "*WWN* [*WWN2*;...]" | **--all**
 Deletes the mapping of one or more device WWNs to a preferred failover *N_Port*. The **--all** option deletes the failover mappings of all device WWNs currently mapped to the specified *N_Port*. The WWN list must be enclosed in double quotation marks. WWNs must be separated by semicolons.
- wwnmappingdisable** "*WWN* [*WWN2*;...]" | **--all**
 Disables one or more device WWN mappings. Use this command if you want to disable the mapping action temporarily without making permanent changes to the mappings. The mappings remain disabled until they are re-enabled or deleted. The **--all** option disables all currently existing device WWN mappings.
- wwnmappingenable** "*WWN* [*WWN2*;...]" | **--all**
 e-enables one or more previously disabled device WWN mappings. The **--all** option re-enables all previously disabled device WWN mappings.
- wwnmapshow**
 Displays all device WWN mappings. For each device WWN, the command displays the *N_Port* number to which it is mapped, the secondary (failover) *N_Port*, and the port group if applicable. The "Current" field shows the port the device is currently using. If the device is not logged in, the field displays "none." If the device is logged in to a port other than the one it is mapped to, the field displays that port. If the device is mapped to a port group, the field displays the number of the port within that port group that the device is currently using. If the device is using a trunk, the field displays which port in that trunk the device is logged in. The "Enabled" field indicates, whether a mapping has been temporarily disabled ("no"), or whether it is in enabled state ("yes").
- reliabilitycounterset** *count*
 Sets the reliability limit for the preferred *N_Port*. This parameter controls the number of ONLINE or OFFLINE State Change Notification (SCN) messages a port can receive before becoming unreliable. The range is 10 through 100. The default value is 25.

The port becomes reliable again if it does not receive any SCN messages for a period of five minutes. Preferred N_Ports settings are not enforced on unreliable N_Ports. When the port becomes reliable again, the behavior is as follows:

- If a FAILBACK flag is set, the port will fail back the configured F_Ports.
- Configured F_Ports that are offline will come back online.
- If LB mode or auto policy is configured, load rebalancing resumes.

--reliabilitycountershow

Displays the configured reliability limit for N_Ports.

--backupmappingsave *N_Port*

Saves the configured F_Ports, static F_Ports for the given N_Port, and F_Ports for which the given N_Port is a preferred one.

--backupmappingdel *N_Port*

Deletes the backup mappings for the given N_Port, if any.

--backupmappingshow *N_Port*

Displays the saved mappings for the given N_Port, if any.

Examples

To display the current state of the Access Gateway with Failover (FO) and Failback (FB) enabled on N_Ports 9 and 12:

```
switch:admin> ag --show
Name                : core_ag
NodeName            : 10:00:00:05:1e:85:ae:f8
Number of Ports    : 40
IP Address(es)     : 10.17.31.2
Firmware Version   : v7.1.0ING
N_Ports            : 8
F_Ports            : 5
Policies enabled   : pg
Persistent ALPA    : Disabled
Static WWN Map     : None
Port Group information :
  PG_ID  PG_Members          PG_Name  PG_Mode
  -----
  0      0;1;2;3;8;          pg0      -
         9;10;11;32;33;
         34;35;36;37;38;
         39
  -----
Fabric Information :
  Attached Fabric Name      N_Ports
  -----
  10:00:00:05:1e:d6:f8:c9   0;1;2;3;8;
                             9;10;11
  -----
N_Port information :
  Port  PortID      Attached PWWN      Attached_Switch
  -----
  0     0x010100   50:00:51:ed:6f:8e:60:28   Spirit\
  1     0x010100   50:00:51:ed:6f:8e:60:28   Spirit\
  2     0x010100   50:00:51:ed:6f:8e:60:28   Spirit\
  -----
  F-port\
  \
```

```

3      0x010100  50:00:51:ed:6f:8e:60:28  Spirit\
8      0x010d00  20:0d:00:05:1e:d6:f8:c9  Spirit\
9      0x010c00  20:0c:00:05:1e:d6:f8:c9  Spirit\
10     0x010e00  20:0e:00:05:1e:d6:f8:c9  Spirit\
11     0x010f00  20:0f:00:05:1e:d6:f8:c9  Spirit\

```

```

Switch  FO  FB  IP_Addr  F_Ports
-----
0      1  1  10.17.31.170  4;5;23;
0      1  1  10.17.31.170  None
0      1  1  10.17.31.170  None
0      1  1  10.17.31.170  None
13     1  1  10.17.31.170  6;
12     1  1  10.17.31.170  7;
14     1  1  10.17.31.170  None
15     1  1  10.17.31.170  None
-----

```

F_Port information :

```

Port  PortID      Attached PWWN          N_Port\
----- \
4     0x010103  20:00:00:05:1e:85:92:88  0 \
5     0x010102  20:01:00:05:1e:85:92:88  0 \
6     0x010d01  20:02:00:05:1e:85:92:88  8 \
7     0x010c01  20:03:00:05:1e:85:92:88  9 \
23    0x010101  10:00:00:05:1e:65:95:81  0 \

```

Preferred N_port Login Exceeded?

```

-----
None          No
None          No
None          No
None          No
None          No
-----

```

Static N-Port to F-Port Mapping :

```

N-Port      F-Port
-----
0           None
1           None
2           None
3           None
8           None
9           None
10          None
11          None
-----

```

To display the current Access Gateway mode:

```

switch:admin> ag --modeshow
Access Gateway mode is enabled.

```

```

switch:admin> ag --modeshow
Access Gateway mode is NOT enabled.

```

AG group policy commands

To show current policies:

```
switch:admin> ag --policyshow
AG Policy                Policy Name      State
-----
Port Grouping            pg               Enabled
Auto Port Configuration  auto            Disabled
Advanced Device Security ads              Disabled
WWN Based Load Balancing wwnloadbalance  Disabled
-----
```

To enable a port grouping policy:

```
switch:admin> ag --policyenable pg
```

To disable a port grouping policy

```
switch:admin> ag --policydisable pg
```

To enable auto port configuration policy when both policies are disabled and the switch is already disabled:

```
switch:admin> ag --policyenable auto
All Port related Access Gateway configurations will be lost.
Please save the current configuration using configupload.
Do you want to continue? (yes, y, no, n): [no] y
```

To disable auto port configuration policy when the switch is disabled:

```
switch:admin> ag --policydisable auto
Default factory settings will be restored.
Default mappings will come into effect.
Please save the current configuration using configupload.
Do you want to continue? (yes, y, no, n): [no] y
Access Gateway configuration has been restored \
to factory default
```

To enable the ADS policy:

```
switch:admin> ag --policyenable ads
```

To disable the ADS policy:

```
switch:admin> ag --policydisable ads
```

To enable the WWN load balancing policy:

```
switch:admin> ag --policyenable wwnloadbalance
```

To disable the WWN load balancing policy:

```
switch:admin> ag --policydisable wwnloadbalance
```

AG port mapping commands

To display current port mappings and port grouping policies:

```
switch:admin> ag --mapshow
```

| N_Port | Configured _F_Ports | Static _F_Ports | Current _F_Ports | Failover | Failback | PG_ID | PG_Name |
|--------|------------------------|--------------------|---------------------|----------|----------|-------|--------------|
| 0 | 4;5;6 | None | 4;5;6 | 1 | 0 | 2 | SecondFabric |
| 1 | 7;8;9 | None | 7;8;9 | 0 | 1 | 0 | pg0 |
| 2 | 10;11 | None | 10;11 | 1 | 0 | 2 | SecondFabric |
| 3 | 12;13 | None | 12;13 | 0 | 1 | 0 | pg0 |

Explanation of fields in `--mapshow` output:

- *Static F_Ports* are part of static F_Port to N_Port mapping.
- *Current F_Ports* are the F_Ports that are currently online and mapped to a given N_Port either because they are mapped to that N_Port or as a result of N_Port failover.
- *Configured F_Ports* are the F_Ports that are explicitly mapped to this N_Port (saved in config).
- *Failover* and *Failback* indicate whether or not N_Port policy is enabled (1) or disabled (0).
- *PG_ID* is the Port Group ID and *PG_Name* is the Port Group Name.

To clear all F_Ports mapped to the configured primary N_Port 0:

```
switch:admin> ag --mapset 0 ""
F_Port to N_Port mapping has been updated successfully
```

To add F_Ports 4 and 6 to N_Port 0 (observe that Port 0 has no configured F_Ports):

```
switch:admin> ag --mapset 0 "4;6"
F_Port to N_Port mapping has been updated successfully
```

To add F_Port 5 to N_Port 2 (observe that N_Port 2 already has mapped F_Ports):

```
switch:admin> ag --mapadd 2 "5"
```

To display the new mappings:

```
switch:admin> ag --mapshow
```

| N_Port | Configured _F_Ports | Static _F_Ports | Current F_Ports | Failover | Failback | PG_ID | PG_Name |
|--------|------------------------|--------------------|--------------------|----------|----------|-------|--------------|
| 0 | 4;6 | None | 4;6 | 1 | 0 | 2 | SecondFabric |
| 1 | 7;8;9 | None | 7;8;9 | 0 | 1 | 0 | pg0 |
| 2 | 5;10;11 | None | 5;10;11 | 1 | 0 | 2 | SecondFabric |
| 3 | 12;13 | None | 12;13 | 0 | 1 | 0 | pg0 |

To delete F_Port 5 that was mapped to N_Port 2:

```
switch:admin> ag --mapdel 2 "5"
Preferred N_port is set for F_Port[s]
Please delete it before removing primary N_Port
ERROR:Unable to remove F_Port[s] from mapping,
retry the command

switch:admin> ag --prefshow
F_Ports                                Preferred N_Port
-----
10;11                                  0
4;5;6                                  2
7;8;9                                  3
-----

switch:admin> ag --prefdel 5 2

Preferred N_Port is deleted successfully \
for the F_Port[s]

switch:admin> ag --mapdel 2 "5"
F_Port to N_Port mapping has been updated successfully
```

To create and display a static mapping:

```
switch:admin> ag --staticadd 17 1

switch:admin> ag --show

N_Port information :
  Port  PortID      Attached PWWN      Attached_Switch
                                F-port\
-----\
  0   0x010100  50:00:51:ed:6f:8e:60:28  Spirit \
  17  0x010100  50:00:51:ed:6f:8e:60:28  Spirit \

Switch  FO  FB  IP_Addr      F_Ports
-----
  0     1  1   10.17.31.170  4;5;23;
  0     1  1   10.17.31.170  None
-----

F_Port information :
  Port  PortID      Attached PWWN      N_Port Preferred \
-----\
  4   0x010103  20:00:00:05:1e:85:92:88  0     \
  5   0x010102  20:01:00:05:1e:85:92:88  0     \

N_port  Login Exceeded?
-----
None           No
None           No
-----

N_Port Configured Current Failover Failback PG_ID PG_Name
  _F_Ports  _F_Ports
-----
  0     None           None           1           1           N/A           N/A
```

| | | | | | | |
|---------------------------------|---|--------|---|---|-----|-----|
| 17 | 2 | 2 | 0 | 0 | N/A | N/A |
| ----- | | | | | | |
| Static N-Port to F-Port Mapping | | | | | | |
| N-Port | | F-Port | | | | |
| ----- | | | | | | |
| 17 | | 1 | | | | |
| ----- | | | | | | |

AG failover policy commands

To display failover policy settings for all N_Ports:

```
switch:admin> ag --failovershow
N_Port  failover_bit
-----
      0             1
      1             0
      2             1
      3             0
```

To set and display failover and failback policies on a single port:

```
switch:admin> ag --failoverenable 1
Failover policy is enabled for port 1

switch:admin> ag --failoverdisable 0
Failover policy is disabled for port 0

switch:admin> ag --failovershow 0
Failover on N_Port 0 is not supported

switch:admin> ag --failbackdisable 2
Failback policy is disabled for port 2

switch:admin> ag --failbackshow 2
Failback on N_Port 2 is not supported

switch:admin> ag --failbackenable 2
Failback policy is enabled for port 2
```

To display failback policy settings for all the N_Ports:

```
switch:admin> ag --failbackshow
N_Port  failback_bit
-----
      0             0
      1             1
      2             0
      3             1
```

To set and display failback policy settings on a single port:

```
switch:admin> ag --failbackenable 0
Failback policy cannot be enabled since failover
policy is disabled for port 0

switch:admin> ag --failbackenable 2
Failback policy is enabled for port 2
```

```
switch:admin> ag --failbackenable 3
Failback on N_Port 3 is not supported
```

```
switch:admin> ag --failbackenable 2
Failback on N_Port 2 is supported
```

Port Group commands

To display Port Group information:

```
switch:admin> ag --pgshow
PG_ID PG_Name      PG_Mode  N_Ports  F_Ports
-----
0     pg0           lb,mfnm  1;3      10;11
2     SecondFabric  -        0;2      4;5;6
-----
```

To create a port group "FirstFabric" that includes N_Ports 1 and 3 and has login balancing enabled:

```
switch:admin> ag --pgcreate 3 "1;3" -n FirstFabric1 -m "lb"
Port Group 3 created successfully
```

```
switch:admin> ag --pgshow
PG_ID PG_Name      PG_Mode  N_Ports  F_Ports
-----
0     pg0           lb,mfnm  none     none
2     SecondFabric  -        0;2      4;5;6
3     FirstFabric   lb       1;3      10;11
-----
```

To rename the port group with *pgid* 2 to "MyEvenFabric":

```
switch:admin> ag --pgrename 2 MyEvenFabric
Port Group 2 has been renamed as MyEvenFabric successfully
```

```
switch:admin> ag --pgshow
PG_ID PG_Name      PG_Mode  N_Ports  F_Ports
-----
0     pg0           lb,mfnm  none     none
2     MyEvenFabric  -        0;2      4;5;6
3     FirstFabric   lb       1;3      10;11
-----
```

To remove the port group with *pgid* 2:

```
switch:admin> ag --pgremove 2
Port Group 2 has been removed successfully
```

```
switch:admin> ag --pgshow
PG_ID PG_Name      PG_Mode  N_Ports  F_Ports
-----
0     pg0           lb,mfnm  0;2      4;5;6
3     FirstFabric   lb       1;3      10;11
-----
```

To enable managed fabric name monitoring in port group 3:

```
switch:admin> ag --pgsetmodes 3 "mfnm"
Managed Fabric Name Monitoring mode has been \
enabled for Port Group 3
```

```
switch:admin> ag --pgshow
```

| PG_ID | PG_Name | PG_Mode | N_Ports | F_Ports |
|-------|-------------|---------|---------|---------|
| 0 | pg0 | lb,mfnm | 0;2 | 4;5;6 |
| 3 | FirstFabric | lb,mfnm | 1;3 | 10;11 |

To disable managed fabric name monitoring in port group 3:

```
switch:admin> ag --pgdelmodes 3 "mfnm"
Managed Fabric Name Monitoring mode has been disabled for Port Group 3
```

```
switch:admin> ag --pgshow
```

| PG_ID | PG_Name | PG_Mode | N_Ports | F_Ports |
|-------|-------------|---------|---------|---------|
| 0 | pg0 | lb,mfnm | 0;2 | 4;5;6 |
| 3 | FirstFabric | lb | 1;3 | 10;11 |

To get the current fabric name monitoring timeout value:

```
switch:admin> ag --pgfnmtov
Fabric Name Monitoring TOV: 120 seconds
```

To set the fabric name monitoring timeout value to 30 seconds:

```
switch:admin> ag --pgfnmtov 30
```

AG Preferred port information commands

To display preferred port settings for F_Ports:

```
switch:admin> ag --prefshow
```

| F_Ports | Preferred N_Port |
|---------|------------------|
| 10;11 | 0 |
| 12;13 | 1 |
| 4;6 | 2 |
| 7;8;9 | 3 |

To delete secondary port mapping for F_Ports 7, 8 and 9:

```
switch:admin> ag --prefdel "7;8;9" 3
Preferred N_Port is deleted successfully \
for the F_Port[s]
```

To set secondary port mapping for F_Ports 7, 8 and 9:

```
switch:admin> ag --prefset "7;8;9" 3
Preferred N_Port is set successfully \
for the F_Port[s]
```


ADS Policy commands

To set the list of allowed devices for Ports 11 and 12 to 'no access':

```
switch:admin> ag --adsset "11;12" ""
WWN list set successfully as the Allow Lists of \
the F_Port[s]
```

To set the list of allowed devices for Ports 1, 10 and 13 to 'all access':

```
switch:admin> ag --adsset "1;10;13" ""
WWN list set successfully as the Allow Lists of \
the F_Port[s]
```

To remove two devices from the lists of allowed devices for ports 1 and 9:

```
switch:admin> ag --adsdel "3;9" \
"22:03:08:00:88:35:a0:12;22:00:00:e0:8b:88:01:8b"
WWNs removed successfully from Allow Lists of the \
F_Port[s]
```

To add a two new device to the lists of allowed devices for ports 1 and 9:

```
switch:admin> ag --adsadd "3;9" \
"20:03:08:00:88:35:a0:12;21:00:00:e0:8b:88:01:8b"
WWNs added successfully to Allow Lists of the \
F_Port[s]
```

To display the lists of allowed devices on the switch:

```
switch:admin> ag --adsshow
F_Port          WWNs Allowed
-----
1              ALL ACCESS
3              20:03:08:00:88:35:a0:12
               21:00:00:e0:8b:88:01:8b
9              20:03:08:00:88:35:a0:12
               21:00:00:e0:8b:88:01:8b
10             ALL ACCESS
11             NO ACCESS
12             NO ACCESS
13             ALL ACCESS
-----
```

Persistent ALPA configuration commands

To enable persistent ALPA in flexible mode:

```
switch:admin> ag --persistentalpaenable 1 -f
Persistent ALPA mode is enabled
```

To enable persistent ALPA in stringent mode:

```
switch:admin> ag --persistentalpaenable 1 -s
Persistent ALPA mode is enabled
```

To disable persistent ALPA mode:

```
switch:admin> ag --persistentalpaenable 0
Persistent ALPA mode is enabled
```

To display the ALPA database entries for F_Port 5:

```
switch:admin> ag --printalpamap 5
Hash table for Port 5 data
PWWN                ALPA
=====
20:12:00:05:1e:85:92:88      1
20:07:00:05:1e:01:0b:4a      3
```

To attempt to remove a device entry from the database while the device is online and cannot be removed:

```
switch:admin> ag --deletepwwnfromdb \
0:12:00:05:1e:85:92:88 20:08:00:05:1e:01:0b:4a Online. \
Cannot delete an online device
```

To remove a device entry from the database when the device is offline:

```
switch:admin> ag --deletepwwnfromdb \
0:12:00:05:1e:85:92:88
Device 20:13:00:05:1e:85:92:88 successfully deleted
```

To remove a device entry from the database when the device is not present in the table

```
switch:admin> ag --deletepwwnfromdb \
0:12:00:05:1e:85:92:00 20:12:00:05:1e:85:92:00 not found.
Please check the device name
```

To remove the PWWN to ALPA value for port 5 from the database and to verify the removal:

```
switch:admin> ag --clearalpamap 5
ALPA Table for port 5 is cleared

switch:admin> ag --clearalpamap 5

Hash table for Port 5 data
PWWN                ALPA
=====
Hash Table is empty
```

Device WWN mapping commands

To create a WWN to N_Port mapping for two devices.

```
switch:admin> ag --addwwnmapping 8 \
"0:12:00:05:1e:85:92:88; 0:12:00:05:1e:85:92:88"
```

To delete one of the device WWN to N_Port mappings.

```
switch:admin> ag --delwwnmapping 8 \
"0:12:00:05:1e:85:92:88"
```

To create a WWN to port group mapping for all currently mapped devices (this command does not affect devices not already mapped or connecting later).

```
switch:admin> ag --addwwnpgmapping 4 --all
```

To add port 13 as a preferred failover N_Port for a device:

```
switch:admin> ag --addwwnfailovermapping 13 "0:12:00:05:1e:85:92:88"
```

To disable all WWN mappings:

```
switch:admin> ag --wwnmappingdisable --all
```

To display the WWN mappings when WWN load balancing policy is not enabled:

```
switch:admin> ag --wwnmapshow
Static Device Mapping Information:
WWN,          1st N_Port 2nd N_Port PG_ID Current Enabled
-----
25:f7:00:0c:29:00:02:8b 9      None      None      None      yes
25:f7:00:0c:29:00:03:8b 9      None      None      None      yes
25:f7:00:0c:29:00:04:8b 9      None      None      None      yes
25:f7:00:0c:29:00:05:8b 9      None      None      None      yes
25:f7:00:0c:29:00:07:8b 9      None      None      None      yes
25:f7:00:0c:29:00:08:8b 9      None      None      None      yes
25:f7:00:0c:29:00:09:8b 9      None      None      None      yes
25:f7:00:0c:29:00:0a:8b 9      None      None      None      yes
-----

Dynamic Device Mapping Information:
-----
No dynamic mappings in use
-----
```

To display the WWN mappings when WWN load balancing policy is enabled:

```
switch:admin> ag --wwnmapshow
Static Device Mapping Information:
WWN,          1st N_Port 2nd N_Port PG_ID Current Enabled
-----
No static mappings are defined
-----

Dynamic Device Mapping Information:
WWN,          1st N_Port 2nd N_Port PG_ID Current Enabled
-----
10:00:00:06:2b:11:52:df 23      None      0        23       yes
-----
```

To configure and display the reliability counter for the preferred N_Port:

```
switch:admin> ag --reliabilitycounterset 50
```

```
switch:admin> ag --reliabilitycountershow
=====
Reliability Counter = 50
=====
```

Backup mapping commands

To save the configured, static, and preferred mapping of an N_Port:

```
switch:admin> ag --backupmappingsave 43
Configured,static and preferred mappings have been saved \
for the N_port successfully.
```

To display the saved mappings for the given N_Port:

```
switch:admin> ag --backupmappingshow 43
N_Port                                     : 43
Backed-up Configured F_Ports             : 15;16;17
Backed-up Static F_ports                  : 18;19
Backed-up Preferred F_ports              : 27;28
```

To delete the backup mappings for the given N_Port:

```
switch:admin> ag --backupmappingdel 43
Backed up mappings have been deleted for the N_port.
```

SEE ALSO [agAutoMapBalance](#), [portCfgNPort](#), [portCfgNPIVPort](#)

agAutoMapBalance

Controls automatic remapping of F_Ports in AG mode.

Synopsis `agautomapbalance --enable [-fport | -nport]`
`[-pg Port_Group_Number | -all]`

`agautomapbalance --disable [-fport | -nport]`
`[-pg Port_Group_Number | -all]`

`agautomapbalance --force`

`agautomapbalance --show`

`agautomapbalance --help`

Description Use this command to control the automatic rebalancing of F_Ports for login distribution in the event that an F_Port goes offline or an N_Port comes online.

If automatic rebalancing is enabled, and an F_Port goes offline, the remaining F_Port logins are redistributed across the existing N_Ports. Similarly if a new N_Port comes online, some of the F_Port logins being routed through existing N_Ports would be failed over to the new N_Ports. Both operations are potentially disruptive. Disabling automatic rebalancing of login distribution provides a way of avoiding disruptions associated with routine F_Port/N_Port offline/online events.

The default values for **agautomapbalance** are as follows:

- Disable automatic login redistribution when F_Ports go offline.
- Enable automatic login redistribution when N_Ports come online.

Use the **--show** option to display the current configuration of the automatic rebalancing feature. The command output varies depending on current AG policy settings:

- If Port Group Policy is enabled on the switch, the command displays the following information for each configured port group:
 - **PG_ID** - Port Group number
 - **LB mode** - Login Balancing mode: enabled or disabled
 - **nport** - Enabled or disabled
 - **fport** - Enabled or disabled
- If Auto Policy is enabled on the switch, the command displays the status of the automatic rebalancing feature per port type as either disabled or enabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--enable Enables automatic login redistribution upon removal or addition of a specified port type.

- fport**
Enables automatic login redistribution in the event that one or more F_Ports go offline. When automatic login redistribution is enabled, the remaining F_Ports are remapped such that logins are balanced among the existing Imports.
- nport**
Enables automatic login redistribution in the event that one or more N_Ports come online. When automatic login redistribution is enabled, the F_Ports mapped to the current N_Ports are rebalanced among the N_Ports.
- disable**
Disables automatic login redistribution upon removal or addition of a specified port type.
- fport**
Disables automatic login redistribution in the event that one or more F_Ports go offline. When automatic login redistribution is disabled, the remaining F_Ports maintain their existing N_Port mappings.
- nport**
Disables automatic login redistribution in the event that one or more N_Ports come online. When auto map balancing is enabled, the F_Ports mapped to the current N_Ports are rebalanced among the N_Ports.
- pg *Port_Group_Number* | -all**
Specifies the port group number or all port groups. These operands are mutually exclusive and optional with the **--enable** and **--disable** options. When a port group is specified, command execution is targeted to the members of that specific port group. When all port groups are specified, command execution is targeted to all port groups defined in the Access Gateway. The port group options are allowed only when login balancing is enabled on the specified port groups.
- force**
Enforces automatic login redistribution on a one-time basis in the event that automatic login redistribution is disabled for N_Port addition, F_Port removal or both. This command forces rebalancing of the F_Port to N_Port mapping once. It does not affect the configuration settings.
- show**
Displays the auto login distribution configuration.
- help**
Displays the command usage.

Examples

To display the automatic login redistribution settings for port groups 0 and 1:

```
switch:admin> agautomapbalance --show
AG Policy:  pg
-----
PG_ID      LB mode      nport      fport
-----
0          Enabled      Enabled     Disabled
1          Disabled     -          -
-----
```

To display the automatic login redistribution settings for N_Ports and F_Ports.

```
switch:admin> agautomapbalance --show

-----
AG Policy:                               Auto
-----
automapbalance on N_Port Online Event:  Disabled
automapbalance on F_Port Offline Event:  Enabled
-----
```

To disable automatic login redistribution on F_Port offline events:

```
switch:admin> agautomapbalance --disable -fport
```

To enable automatic login redistribution on F_Ports and N_Ports on port group 1 in the Access Gateway:

```
switch:admin> agautomapbalance --enable -fport pg 1
```

To disable automatic login redistribution on F_Ports and N_Ports on all port groups in the Access Gateway:

```
switch:admin> agautomapbalance --disable -all
```

SEE ALSO [ag, agShow](#)

agShow

Displays the Access Gateway information registered with the fabric.

Synopsis

agshow

agshow --name *ag_name*

agshow --local

Description

This command displays the details of the F_Ports and the configured N_Ports in the Access Gateway attached to the fabric. The command output displays the following information.

Name

The name of the Access Gateway.

NodeName

The World Wide Name of the Access Gateway node.

Ports

The number of ports in the Access Gateway.

Enet IP Addr

The IP address of the Access Gateway.

Firmware

Current firmware running on the Access Gateway.

Local/Remote

Indicates whether the Access Gateway is locally or remotely registered to this switch.

World Wide Name

The world wide name (WWN) of the given Access Gateway.

N-Port ID(s)

The port ids of the N_Ports configured in the given Access Gateway.

N-Ports

The number of configured N_Ports that are online.

F-Ports

The number of F_Ports that are online.

Attached F-Port information

Displays the Port ID, the switch F-port number, and the Port WWN of each F_Port that is online on the Access Gateway.

Access Gateway F-Port information

Displays the port number, the port ID, and the Port WWN of the Access Gateway to which the F_Port is connected. This information is displayed only if both the Access Gateway and the switch to which the Access Gateway is attached are running Fabric OS v.7.0.0 or later. The data in this section is not always synchronized with the data in the Attached F-Port information section.

Notes NPIV capability should be enabled on the ports connected to the Access Gateway. NPIV capability is enabled by default. Use **portCfgNPIVPort** to enable NPIV capability a port if it was previously disabled. Note that enabling or disabling NPIV capability is no longer supported in Access Gateway mode.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

--name *ag_name* Displays information regarding a specific Access Gateway that is registered with this fabric.

--local Display information regarding all Access Gateways that are locally registered to this switch.

Examples To display the Access Gateway information registered with the fabric:

```
switch:admin> agshow --name core_ag
Name                : core_ag
NodeName            : 10:00:00:05:1e:85:ae:f8
N-Port ID(s)       : 0x010100,0x010d00,0x010c00,0x010e00,0x010f00
Number of Ports    : 40
IP Address(es)     : 10.17.31.2
Firmware Version   : v7.1.0ING
N-Ports            : 5
F-Ports            : 5
Attached F-Port information :
PortID      Port WWN                               Switch F-port
-----
0x010101    10:00:00:05:1e:65:95:81                        0
0x010102    20:01:00:05:1e:85:92:88                        0
0x010d01    20:02:00:05:1e:85:92:88                        13
0x010c01    20:03:00:05:1e:85:92:88                        12
Access Gateway F-Port Information :
F-Port NumberF-Port IDF-Port WWN
-----
23      0x01010120:17:00:05:1e:85:ae:f8
```

To display the locally registered Access Gateways:

```
switch:admin> agshow --local
Worldwide Name      Ports Enet IP Addr Firmware \
-----\
10:00:00:05:1e:04:06:ae  24  10.32.173.64  v6.4.0  \

Local/Remote Name
-----
local      L5D_B14_4024_1
```

To display all Access Gateways attached to the fabric:

```
switch:admin> agshow
Worldwide Name      Ports Enet IP Addr Firmware \
-----\
10:00:00:05:1e:02:b7:2c  16  10.32.173.62  v6.4.0  \
10:00:00:05:1e:04:06:ae  24  10.32.173.64  v6.4.0  \
10:00:00:05:1e:04:06:ae  24  10.32.173.51  v6.4.0  \
```

agShow

```
Local/Remote Name
-----
local          L5D_B10_4016_1
local          L5D_B14_4024_1
local          L5D_B13_200_AG
```

SEE ALSO [portCfgNPIVPort](#)

aliAdd

Adds a member to a zone alias.

Synopsis `aliadd "aliName", "member[; member...]"`

Description Use this command to add one or more members to an existing zone alias. The alias member list cannot contain another zone alias.

This command changes the defined configuration. For the change to become effective, enable the zone configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the zone configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"aliName"

Specify the name of a zone alias, enclosed in double quotation marks.

"member"

Specify a member or list of members to be added to the alias, enclosed in double quotation marks. Members must be separated by semicolons. An alias member can be specified by one or more of the following methods:

- A switch domain and port index pair. Use **switchShow** for a list of valid port index numbers.
- A world wide name (WWN).

Examples To add members to zone aliases array1, array2, and loop1:

```
switch:admin> aliadd "array1", "1,2"
```

```
switch:admin> aliadd "array2", "21:00:00:20:37:0c:72:51"
```

SEE ALSO [aliCreate](#), [aliDelete](#), [aliRemove](#), [aliShow](#)

aliCreate

Creates a zone alias.

Synopsis `alicreate "aliName", "member[: member...]"`

Description Use this command to create a new zone alias. The zone alias member list must have at least one member (empty lists are not allowed). The alias member list cannot contain another zone alias. Refer to the **zoneCreate** command for more information on name and member specifications.

This command changes the defined configuration. For the change to become effective, enable the zone configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the zone configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"aliName"

Specify a name for the zone alias, in double quotation marks. A zone alias name must begin with a letter and can be followed by any number of letters, numbers, and underscore characters. Names are case-sensitive. For example, "Ali_1" and "ali_1" are different zone aliases. Spaces are ignored.

"member"

Specify a member or list of members to be added to the alias, enclosed in double quotation marks. Members must be separated by semicolons. An alias member can be specified by one or more of the following methods:

- A switch domain and port index pair. Use **switchShow** for a list of valid port index numbers.
- A world wide name (WWN).

Examples To create a zone alias defined by domain and port index pairs:

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34"
```

To create a zone alias with one member defined by WWN.

```
switch:admin> alicreate "array2", "21:00:00:20:37:0c:66:23"
```

SEE ALSO [aliAdd](#), [aliDelete](#), [aliRemove](#), [aliShow](#)

aliDelete

Deletes a zone alias.

Synopsis `alidelete "aliName"`

Description Use this command to delete a zone alias.

This command changes the defined configuration. For the change to become effective, enable the zone configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the zone configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operand is required:

"aliName"

Specify the name of the zone alias to be deleted. Double quotation marks are optional.

Examples To delete the zone alias "array2":

```
switch:admin> alidelete "array2"
```

```
switch:admin> alidelete array1
```

SEE ALSO [aliAdd](#), [aliCreate](#), [aliRemove](#), [aliShow](#)

aliRemove

Removes a member from a zone alias.

Synopsis `aliremove "aliName", "member[; member...]"`

Description Use this command to remove one or more members from an existing zone alias.

If all members are removed, the zone alias is deleted.

This command changes the defined configuration. For the change to become effective, enable the zone configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the zone configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands This command has the following operands:

"aliName"

Specify the name of the zone alias from which members are to be removed in double quotation marks. This operand is required.

"member"

Specify a member or list of members to be removed from the alias. The list must be enclosed in double quotation marks. Members must be separated by semicolons. An alias member can be specified by one or more of the following methods:

- A switch domain and port index number pair. Use **switchShow** for a list of valid port index numbers.
- A world wide name (WWN)

The member list is located by an exact string match; therefore, it is important to maintain the order when removing multiple members. For example, if a zone alias contains "1,2; 1,3; 1,4", then removing "1,3; 1,4" succeeds but removing "1,4; 1,3" fails.

Examples To remove a world wide name from "array1":

```
switch:admin> aliremove "array1", "3,5"
```

```
switch:admin> aliremove "array1", "21:00:00:20:37:0c:76:8c"
```

```
switch:admin> aliremove "array1", "OxEF"
```

SEE ALSO [aliAdd](#), [aliCreate](#), [aliDelete](#), [aliShow](#)

aliShow

Displays zone alias information.

Synopsis `alishow [--ic] ["pattern"][, mode]`

`alishow --help`

Description Use this command to display zone configuration information.

Use the pattern operand to display only matching zone alias names in the defined configuration.

If no parameters are specified, all zone configuration information (both defined and effective) is displayed. Refer to **cfgShow** for a description of this display.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands When invoked without an operand, this command displays all zone configuration information (defined and effective). Refer to **cfgShow** for a description of this display. The following operands are optional:

`--ic`

Displays all configured zone alias names for a given pattern without case distinction.

`"pattern"`

A POSIX-style regular expression that matches zone alias names. This operand must be enclosed in quotation marks. Patterns may contain:

- Question mark (?) - matches any single character.
- Asterisk (*) - matches any string of characters.
- Range - matches any character within the range. Ranges must be enclosed in brackets: for example, [0-9] or [a-f].

`mode`

Specify 0 to display the contents of the transaction buffer (the contents of the current transaction), or specify 1 to display the contents of the nonvolatile memory. The default value is 0.

`--help`

Displays the command usage.

Examples To display all zone aliases beginning with "arr":

```
switch:admin> alishow "arr*"
alias: array1  21:00:00:20:37:0c:76:8c
alias: array2  21:00:00:20:37:0c:66:23
```

aliShow

To display all zone alias names beginning with "arr", regardless of the case:

```
switch:admin> alishow --ic "arr"  
alias: array1 20:e0:00:05:33:11:1f:00  
alias: ARRAY2 2f:11:00:05:33:c1:37:a2
```

SEE ALSO [aliAdd](#), [aliCreate](#), [aliDelete](#), [aliRemove](#)

appLoginHistory

Displays the history of HTTP login sessions.

Synopsis `apploginhistory --show`

`apploginhistory --help`

Description Use this command to display the history of HTTP login sessions from external management applications such as Brocade Network Advisor or Web Tools. The command displays both current sessions and a history of past sessions. For each entry, the command output shows the following information:

- The date and time when the session started (YYYY/MM/DD-HH:MM:SS.MS).
- The IP address of the machine that initiated the HTTP login.
- The role of the user initiating the login.
- The application initiating the login.

The history supports a maximum of 100 entries. If the history exceeds the maximum size, the oldest entries will be removed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--show Displays the currently active logins and a history of past logins.

--help Displays the command usage.

Examples To display the application login history:

```
switch:admin> apploginhistory -- show
```

```
The following are history of logins to the switch from external
applications
```

```
(Login Timestamp, IP Address, User Account, Application):
```

```
2011/10/13-15:51:31.203814, 10.20.10.160, admin, DCFM-HttpConnector/null
2011/10/13-15:51:33.095638, 10.20.10.160, admin, DCFM-HttpConnector/null
2011/10/13-15:52:14.178299, 10.20.10.161, admin, DCFM-HttpConnector/null
2011/10/13-15:52:14.643639, 10.20.10.161, admin, DCFM-HttpConnector/null
2011/10/13-15:53:55.966050, 10.20.10.161, admin, DCFM-HttpConnector/null
2011/10/13-15:53:56.329979, 10.20.10.161, admin, DCFM-HttpConnector/null
2011/10/13-15:53:56.858216, 10.20.10.161, admin, DCFM-HttpConnector/null
2011/10/13-15:53:57.547672, 10.20.10.161, admin, DCFM-HttpConnector/null
```

```
The following are the sessions from the external applications that are
```

appLoginHistory

active currently:

```
2011/10/13-15:51:31.203814, 10.20.10.160, admin, DCFM-HttpConnector/null
2011/10/13-15:51:33.095638, 10.20.10.160, admin, DCFM-HttpConnector/null
2011/10/13-15:52:14.178299, 10.20.10.161, admin, DCFM-HttpConnector/null
2011/10/13-15:53:57.547672, 10.20.10.161, admin, DCFM-HttpConnector/null
```

SEE ALSO None

aptPolicy

Changes or displays the Advanced Performance Tuning (APT) policy.

Synopsis `aptpolicy [policy]`

Description Use this command to display and change the advanced performance tuning (APT) policies on a switch.

Distributed path selection (DPS) is supported in logical fabrics. APT policy settings affecting the DPS behavior can be configured per logical switch, and settings apply to the partition for which they are set. Note that policy settings for the base switch or any switch in the base fabric affect all traffic going through the base fabric including any logical fabric traffic that uses the base fabric.

When invoked without arguments, this command displays the APT policies supported on this switch, as well as the current policy.

Notes You must disable the switch before using this command to change the current policy. Changes take effect immediately for all EX/VEX_Ports after the switch is re-enabled.

For details on performance tuning, refer to the *Fabric OS Administrator's Guide*.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

policy

Specifies the APT policy. The following policies are supported:

1

Port-based routing policy. With this policy, the path chosen for an ingress frame is based on both of the following items:

- The ingress port on which the frame was received.
- The destination domain for the frame.

The chosen path remains the same if Dynamic Load Sharing (DLS) is not enabled. If DLS is enabled, a different path may be chosen for a fabric event. Refer to **disSet** for a definition of a fabric event.

This policy may provide better ISL utilization when there is little or no oversubscription of the ISLs.

Note that static routes are supported only with this policy.

2

Device Based Routing Policy. Device-based routing is supported in FICON environments, and in open environments only when FICON coexists. With this policy, the path chosen for an ingress frame is based on the following items:

- The ingress port on which the frame was received.
- The FC address of the source fabric device (SID) for this frame.
- The FC address of the destination fabric device (DID) for this frame.

This policy optimizes the utilization of the available paths by allowing I/O traffic between different source fabric device (SID) or destination fabric device (DID) pairs to use different paths. As a result, every distinct flow in the fabric can take a different path through the fabric. Effectively, device based routing works the same as exchange-based routing but does not use the Originator Exchange ID (OXID) field. This helps to ensure that the exchanges between a pair of devices stay in order.

Device-based routing is also a form of Dynamic Path Selection (DPS). DPS assigns communication paths between end devices in a fabric to egress ports in ratios proportional to the potential bandwidth of the ISL, ICL, or trunk group. When there are multiple paths to a destination, the input traffic is distributed across the different paths in proportion to the bandwidth available on each of the paths. This improves utilization of the available paths and reduces possible path congestion.

3

Exchange-based routing policy (default). With this policy, the path chosen for an ingress frame is based on all of the following items:

- The ingress port on which the frame was received.
- The FC address of the SID for this frame.
- The FC address of the DID for this frame.
- The FC OXID for this frame.

This policy optimizes the utilization of the available paths by allowing I/O traffic between different SID, DID, or OXID pairs to use different paths. All frames received on an ingress port with the same SID, DID, or OXID parameters take the same path unless there is a fabric event. Refer to **dlsSet** for the definition of a fabric event.

This policy does not support static routes. DLS is always enabled and the DLS setting cannot change with this policy.

Examples To display the current APT policy:

```
switch:admin> aptpolicy
      Current Policy: 3
3: Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy (FICON support only)
3: Exchange Based Routing Policy
```

To change the current APT policy to the exchange-based routing policy:

```
switch:admin> aptpolicy 3
Switch must be disabled in order to modify \
this configuration parameter. To disable the switch, \
use the "switchDisable" command.

switch:admin> switchdisable

switch:admin> aptpolicy 3
Policy updated successfully.

switch:admin> switchenable

switch:admin> aptpolicy
Current Policy: 3

3: Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy (FICON support only)
3: Exchange Based Routing Policy
```

SEE ALSO [dlsReset](#), [dlsSet](#), [dlsShow](#), [switchDisable](#)

auditCfg

Modifies and displays the audit log filter configuration.

Synopsis `auditcfg --class audit_class`
`auditcfg --enable | --disable`
`auditcfg --severity severity_level`
`auditcfg --show`

Description Use this command to configure the audit logging and to display the audit log configuration. This command allows you to set filters by configuring certain classes, to add or remove any of the classes in the filter list, to set severity levels for audit messages, and to enable or disable audit filters. Based on the configuration, certain classes are logged to syslog for auditing. Syslog configuration is required for logging audit messages. Use the **syslogAdmin** command to add the syslogd server IP address.

Notes The RAS class is not configurable; its function is to audit the audit log management operations and it is always enabled internally.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--class Configures filters for a specified audit class. To add or remove any of the classes in the filter list, reissue the **--class** option.

audit_class Specifies the filters to be configured. Valid values are: 1-ZONE, 2-SECURITY, 3-CONFIGURATION, 4-FIRMWARE, 5-FABRIC, 7-LS (Logical Switch), 8-CLI, and 9-MAPS. The filter is specified by its numeric value. To add more than one filter, the numeric values must be separated by commas. Spaces are not permitted. This operand is required.

--severity Sets audit severity level to a specified value. When severity is set, only log messages of type *severity_level* and higher are displayed. You cannot enter multiple severity levels.

severity_level Valid values are INFO, WARNING, ERROR, and CRITICAL. By default, all messages are logged. This operand is required.

--enable Enables all filters. This action enables an existing configuration; it does not change the configuration.

--disable Disables all filters. This action disables an existing configuration; it does not change the configuration.

--show

Displays the current configuration. This operand is optional.

Examples

To configure the audit log filter:

```
switch:admin> auditcfg --class 2,3,8
Audit filter is configured.
```

To enable audit logging:

```
switch:admin> auditcfg --enable
Audit filter is enabled.
```

To display the configuration:

```
switch:admin> auditcfg --show
Audit filter is enabled.
2-SECURITY
3-CONFIGURATION
8-CLI
Severity level: INFO
```

To disable audit logging and to display the configuration (the filters are unchanged but show the disabled state):

```
switch:admin> auditcfg --disable
Audit filter is disabled.
switch:admin> auditcfg --show
Audit filter is disabled.
2-SECURITY
3-CONFIGURATION
8-CLI
Severity level: INFO
```

SEE ALSO

[auditDump](#), [rasAdmin](#)

auditDump

Displays or clears the audit log.

Synopsis `auditdump -s | -show`

`auditdump -c | -clear`

Description Use this command to display or clear the audit log on the switch. The audit log persistently saves the most recent 256 log entries on the switch. On modular platforms, the entries are not shared across CPs. Each CLI executed on a switch from non-root user account is saved as an audit log. To display or clear the logs, this command must be issued for each CP separately.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-s | -show

Displays the audit log on the switch or the CP.

-c | -clear

Clears the audit log on the switch or the CP.

Examples To display the audit log:

```
switch:admin> auditdump -s
0 AUDIT, 2010/02/18-01:45:52 (PDT), [SEC-3022], INFO, SECURITY,
admin/admin/10.32.220.161/http/DCFM-HttpConnec, ad_0/ras035/FID
128,, Event: logout, Status: success, Info: Successful logout
by user [admin].

1 AUDIT, 2010/02/18-01:46:00 (PDT), [SEC-3022], INFO, SECURITY,
admin/admin/10.32.220.161/http/DCFM-HttpConnec, ad_0/ras035/FID
128,, Event: logout, Status: success, Info: Successful logout
by user [admin].

2 AUDIT, 2010/02/18-01:46:18 (PDT), [SEC-3020], INFO, SECURITY,
admin/admin/10.32.220.161/http/DCFM-HttpConnec, ad_0/ras035/FID
128,, Event: login, Status: success, Info: Successful login
attempt via HTTP, IP Addr: 10.32.220.161.

3 AUDIT,2012/05/23-03:45:15 (UTC),[RAS-3005], INFO, CLI,
admin/admin/NONE/console/CLI,ad_0/McKsSpirit/CHASSIS,
CLI:clihistory

4 AUDIT,2012/05/23-04:12:04 (UTC),[RAS-3005],INFO,CLI,
admin/admin/NONE/console/CLI,ad_0/McKsSpirit/CHASSIS,
CLI:auditdump -s

(output truncated)
```


To clear the audit log:

```
switch:admin> auditdump -c
```

SEE ALSO [auditCfg](#)

authUtil

Displays and sets the authentication configuration.

Synopsis

authutil

authutil --show

authutil --set *option value*

authutil --policy -sw *option* | -dev *option*

authutil --authinit [*slot*]*port*[, [*slot*]*port*...] | **allE**

Description

Use this command to display and set local switch authentication parameters.

Use **--set** to change authentication parameters such as protocol, Diffie-Hellman group (DH group), or hash type. When no protocol is set, the default setting of "FCAP, DH- CHAP" is used. When no group is set, the default setting of "" (meaning "0,1,2,3,4") is used. Configuration settings are saved persistently across reboots. Configuration changes take effect during the next authentication request.

Use the **--show** command to display the current authentication configuration.

Authentication parameters are set on a per-switch basis. If Virtual Fabrics are enabled, all authentication parameters apply to the current logical switch context only, and must be configured separately for each logical switch. Use **setContext** to change the current logical switch context.

In a VF environment, authentication is performed only on physical E_Ports, not on logical interswitch links (LISLs).

An FCR switch does not depend on the authentication policy to perform authentication or encryption/compression key exchange with the edge fabric. The authentication policy set on an FCR switch is not considered to perform authentication with the edge fabric. The **authUtil** configuration on the FCR switch does not affect EX_Port. The EX_Port acts passively by accepting the parameters received from the edge fabric.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

--show

Displays the local authentication configuration. This option is supported in the Access Gateway (AG) mode.

--set option value

Modifies the authentication configuration. Valid *options* and their *values* include the following:

-a fcap | dhchap | all

Sets the authentication protocol. Specify "fcap" to set only FCAP authentication. Specify "dhchap" to set only DH-CHAP authentication. Specify "all" to set both FCAP and DH-CHAP, which is the default setting. When authentication is set to "all", the implicit order is FCAP followed by DH-CHAP. This means that in authentication negotiation, FCAP is given priority over DH-CHAP on the local switch. If the negotiation is done for an encrypted port, DHCHAP takes precedence over FCAP. The **--set dhchap** and **--set all** options are supported in the AG mode.

Before setting the authentication type, make sure that FCAP certificates and dhchap secrets are configured.

-g 0 | 1 | 2 | 3 | 4 | *

Sets the Diffie-Hellman (DH) group. Valid values are 0 to 4 and "*". The DH group 0 is called NULL DH. Each DH group implicitly specifies a key size and associated parameters. A higher group value provides stronger cryptography and a higher level of security. When DH group is set to a specified value, only that DH group is enabled. Specifying "*" enables all DH groups 0, 1, 2, 3, and 4, in that order. This means that in authentication negotiation, the NULL DH group is given priority over all other groups. In the case of a port that is enabled for encryption and you specify "*", the DH group 4 is selected. This option is supported in AG mode.

-h sha256 | sha1 | md5 | all

Sets the hash type. Valid values are "sha256", "sha1", "md5", or "all". Enabling sha256 access is required when configuring the system for FIPS. Refer to the *Fabric OS Administrator's Guide* for details on FIPS configuration.

--policy

Sets the switch authentication policy or device authentication policy. The following options are supported:

-sw off | passive | active | on

Sets the switch authentication policy. Specify one of the following modes. Operands are exclusive. If the switch has ports enabled for encryption, only the ON and ACTIVE options are supported. Only **on** and **off** options are supported in AG mode.

off

Turns the authentication policy off, and the switch rejects any authentication requests.

passive

Sets the authentication policy to passive mode (default). The switch does not initiate authentication but participates in authentication if the connecting switch initiates authentication.

active

Sets the authentication policy to active mode. During switch initialization, authentication is initiated on all E_Ports, but the port is not disabled if the connecting switch does not support authentication or the authentication policy is turned off.

- on**
Sets the switch authentication policy to ON mode. Strict authentication is enforced on all E_Ports. The interswitch link (ISL) goes down (port disable), if the connecting switch does not support the authentication or the authentication policy is switched off.
- dev off | passive | on**
Sets the device authentication policy. Three modes are supported. Device authentication policy is off by default. This option and suboptions are supported in AG mode.
- off**
Turns off the device authentication policy. Authentication is not required. The switch ignores any authentication requests and continues with the FC probing without authentication.
- passive**
Sets the authentication policy to passive mode. Authentication is optional. If the attached device is capable of doing the authentication then the switch participates in authentication; otherwise it forms an F_Port without authentication. In this mode the device accepts authentication on all F_Ports.
- on**
Sets the authentication policy to "on" mode. Authentication is mandatory. If the attached device is not capable of doing authentication, the corresponding port is disabled.
- authinit [slot]port [, [slot]port...] allE**
Reinitiates authentication on selected ports after changing the DH-CHAP group, hash type, and shared secret between a pair of switches. This command does not work on Private, Loop, NPIV and FICON devices. This command may bring down the E_Ports if the DH-CHAP shared secrets are not installed correctly. This command is not supported on encrypted ports. This option is not supported in AG mode. This command does not initiate authentication on disabled ports. Valid options include the following:
- slot**
Specify the slot number, if applicable, followed by a slash (/).
- port**
Specify the port number. On enterprise-class platforms, use the *slot/port* format for specifying the port number.
- allE**
Specify all E_Ports in the switch.

Examples

To set DH-CHAP as the authentication protocol:

```
switch:admin> authutil --set -a dhchap
Authentication is set to dhchap.
```

To set both protocols in order of FCAP and then DH-CHAP:

```
switch:admin> authutil --set -a all
Authentication is set to fcap,dhchap.
```

To enable sha1 hash type:

```
switch:admin> authutil --set -h sha1
Hash is set to sha1.
```

To enable all hash types:

```
switch:admin> authutil --set -h all
Hash is set to sha256,sha1,md5.
```

To check the hash types set:

```
switch:admin> authutil --show
AUTH TYPE      HASH TYPE      GROUP TYPE
-----
fcap,dhchap    sha256,sha1,md5    1

Switch Authentication Policy: PASSIVE
Device Authentication Policy: OFF
```

To enable Sha256 hash type:

sha256 hash type is recommended for FIPS configuration.

```
switch:admin> authutil --set -h sha256
Hash is set to sha256.
```

```
switch:admin> authutil --show
AUTH TYPE      HASH TYPE      GROUP TYPE
-----
fcap,dhchap    sha256          1

Switch Authentication Policy: PASSIVE
Device Authentication Policy: OFF
```

To set DH group 3:

```
switch:admin> authutil --set -g 3
DH Group was set to 3.
```

To set all DH groups to be specified in the authentication negotiation in the order of 0, 1, 2, 3, and 4:

```
switch:admin> authutil --set -g ""
DH Group is set to 0,1,2,3,4
```

To set the Switch policy to active mode:

```
switch:admin> authutil --policy -sw active
Warning: Activating the authentication policy requires
either DH-CHAP secrets or PKI certificates depending
on the protocol selected. Otherwise, ISLs will be
segmented during next E-port bring-up.
ARE YOU SURE (yes, y, no, n): [no] y
Auth Policy is set to ACTIVE
```

To set the Device policy to passive mode:

```
switch:admin> authutil --policy -dev passive
Warning: Activating the authentication policy requires
DH-CHAP secrets on both switch and device. Otherwise,
the F-port will be disabled during next F-port
bring-up.
ARE YOU SURE (yes, y, no, n): [no] y
Device authentication is set to PASSIVE
```

To set the device authentication policy to "on" mode:

```
switch:admin> authutil --policy -dev on
Warning: Activating the authentication policy requires
DH-CHAP secrets on both switch and device. Otherwise,
the F-port will be disabled during next F-port
bring-up.
ARE YOU SURE (yes, y, no, n): [no] y
Device authentication is set to ON
2008/03/24-23:13:06, [AUTH-1003], 112,, INFO, Stealth_3,
Device authentication type has been successfully set to ON
```

To start authentication on E/F_Ports 2, 3, and 4:

```
switch:admin> authutil --authinit 2,3,4
Warning: Initiating the authentication requires either DH-CHAP
secrets or PKI certificates depending on the protocol selected.
Failed authentication may result in traffic disruption.
Authentication will not be initiated on encrypted ports.
ARE YOU SURE (yes, y, no, n): [no]y
```

SEE ALSO [portShow](#), [secAuthSecret](#)

bannerSet

Sets the banner on the local switch.

Synopsis `bannerSet [banner]`

Description Use this command to set the banner on the local switch.

The banner is a string of alphanumeric characters. It is displayed after you log in to a switch.

The banner can be created using the *banner* operand or interactively by entering the **bannerSet** command without an operand.

If you enter the banner text using the interactive method, the valid length is 1022 characters. If the banner text length exceeds the maximum allowed, the software truncates the input. To close the banner text string, enter a period at the beginning of a new line.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operand:

banner

Specify a text string to be displayed upon login. If you enter the banner text using the banner operand, the valid length is 116 characters.

Examples To set a new banner for a switch:

```
switch:admin> bannerSet "My banner"
```

```
switch:admin> bannerSet
```

```
Please input context of security banner (press "." RETURN \
at the beginning of a newline to finish input):
```

SEE ALSO [bannerShow](#), [motd](#)

bannerShow

bannerShow

Displays the banner text.

Synopsis **bannershow**

Description Use this command to display the text of the local switch banner.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the banner for a switch:

```
switch:admin> bannershow  
Banner: Do not disturb the setup on this switch.
```

SEE ALSO [bannerSet](#), [motd](#)

bcastShow

Displays broadcast routing information.

Synopsis `bcastshow`

Description Use this command to display the broadcast routing information for all ports in the switch. The broadcast routing information indicates all ports that are members of the broadcast distribution tree: ports that are able to send and receive broadcast frames.

Normally, all F_Ports and FL_Ports are members of the broadcast distribution tree. The broadcast path selection protocol selects the E_Port members of this tree in a manner designed to prevent broadcast routing loops.

The following fields are displayed:

Group The multicast group ID of the broadcast group (always 256).

Member Fx_Ports A map of all F_Ports and FL_Ports ports in the broadcast tree.

Member E_Ports A map of all E_Ports in the broadcast tree.

The broadcast routing information for the ports is displayed as a set of hexadecimal bit maps. Each bit in a bit map represents a port, with the least significant bit in each row representing port 0, 32, 64, and so on.

Notes The output from this command may vary depending on the hardware platform.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the broadcast routing information for all ports in the switch:

```
switch:admin> bcastShow

Group  Member Ports      Member ISL Ports
-----
256    0x00012083        0x00002080
       0x00000440        0x00000400
       0x00770000        0x00700000
       0x00008200        0x00000000
       0x00000001        0x00000000
```

In this example from a switch with 128 ports, the member ports consist of ports 7, 13, 42, 84, 85, and 86. The final Member Ports bit set represents the embedded port (frames sent to be handled by firmware) and is typically set.

SEE ALSO [portRouteShow](#)

bladeCfgGeMode

Configures the operating mode of the Brocade FX8-24 blade to allow 1 GbE, 10 GbE, or both.

Synopsis `bladecfggemode --set mode -slot slot`

`bladecfggemode --show -slot slot | -all`

`bladecfggemode --help`

Description Use this command to configure the GbE port mode on the Brocade FX8-24 extension blade or to display the configuration. The mode configuration controls which ports are enabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported only on the Brocade FX8-24 blade in a DCX 8510 chassis.

Operands This command has the following operands:

--set mode

Sets the GbE port mode for a specified slot. Valid modes are one of the following:

1g

Enables the GbE ports ge0-ge9 (xge0 and xge1 are disabled).

10g

Enables the GbE ports xge0 and xge1 (ge0-ge9 ports are disabled).

dual

Enables the GbE ports ge0-ge9 and xge0 (xge1 is disabled).

-slot slot

Specifies the slot number for the FX8-24 blade. This operand is required when setting the GbE port mode

--show

Displays the GbE port mode for the specified slots.

-slot slot

Displays the GbE port mode for a single slot.

-all

Displays the GbE port mode for all configured slots.

--help

Displays the command usage.

Examples To configure the Brocade FX8-24 blade in slot 4 in 1G mode:

```
switch:admin> bladecfggemode --set 1g -slot 4
```

To display the GbE port mode for the Brocade FX8-24 blade in slot 4:

```
switch:admin> bladecfggemode --show -slot 4  
bladeCfgGeMode: Blade in slot 4 is configured in 10GigE Mode  
10GigE mode: only xge0 and xge1 are enabled (ge0-9 ports \  
are disabled)
```

To display the GbE port mode for all configured slots:

```
switch:admin> bladecfggemode --show -all  
bladeCfgGeMode: Blade in slot 1 is configured in 1GigE Mode  
1GigE mode: only the ge0-9 ports are enabled (xge0 and xge1 \  
are disabled)  
bladeCfgGeMode: Blade in slot 4 is configured in 10GigE Mode  
10GigE mode: only xge0 and xge1 are enabled (ge0-9 ports \  
are disabled)
```

SEE ALSO None

bladeDisable

Disables all user ports on a blade.

Synopsis `bladedisable slot`

Description Use this command to disable all user ports on a blade. All Fibre Channel ports on the blade are taken offline. If the switch was connected to a fabric through this blade, the remaining switches reconfigure, and the switch reconfigures based on the other blade ports. As each port is disabled, the front panel LED changes to a slow-flashing amber.

After issuing **bladeDisable** on a slot in a chassis, **switchShow** displays the user ports in the disabled state. The blade is still shown as enabled in both **switchShow** and **slotShow** output.

The blade must be disabled before making configuration changes or before running many of the diagnostic tests. The blade does not need to be disabled before rebooting or powering off.

You cannot disable a blade when the blade is faulted, powered off, or running diagnostics.

This command disables the ports on a single blade. To disable the ports in an entire chassis, use the **chassisDisable** command.

Notes This command appears to execute successfully along with **slotpoweron** operation on an AP Blade, but the command is ignored.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

slot
Specifies the slot number for the ports to be disabled.

Examples To disable a blade in slot 5:

```
switch:admin> bladedisable 5
Blade 5 is being disabled...Done
```

To verify that the user ports on the blade are disabled:

```
switch:admin> switchshow -slot 5
switchName:      DCX57_9
switchType:      62.1
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     9
switchId:        fffc09
switchWwn:       10:00:00:05:1e:40:4a:00
zoning:          ON (all_cfg)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
```

| Slot | Blade Type | ID | Model Name | Status |
|------|------------|----|------------|---------|
| 5 | COREBLADE | 52 | CORE8 | ENABLED |

| Index | Slot | Port | Address | Media | Speed | State | Proto |
|-------|------|------|---------|-------|-------|---------|-------------|
| 384 | 5 | 0 | ----- | cu | 8G | In_Sync | FC Disabled |
| 385 | 5 | 1 | ----- | cu | 8G | In_Sync | FC Disabled |
| 386 | 5 | 2 | ----- | cu | 8G | In_Sync | FC Disabled |
| 387 | 5 | 3 | ----- | cu | 8G | In_Sync | FC Disabled |
| 388 | 5 | 4 | ----- | cu | 8G | In_Sync | FC Disabled |
| 389 | 5 | 5 | ----- | cu | 8G | In_Sync | FC Disabled |
| 390 | 5 | 6 | ----- | cu | 8G | In_Sync | FC Disabled |
| 391 | 5 | 7 | ----- | cu | 8G | In_Sync | FC Disabled |
| 392 | 5 | 8 | ----- | cu | 8G | In_Sync | FC Disabled |
| 393 | 5 | 9 | ----- | cu | 8G | In_Sync | FC Disabled |
| 394 | 5 | 10 | ----- | cu | 8G | In_Sync | FC Disabled |
| 395 | 5 | 11 | ----- | cu | 8G | In_Sync | FC Disabled |
| 396 | 5 | 12 | ----- | cu | 8G | In_Sync | FC Disabled |
| 397 | 5 | 13 | ----- | cu | 8G | In_Sync | FC Disabled |
| 398 | 5 | 14 | ----- | cu | 8G | In_Sync | FC Disabled |
| 399 | 5 | 15 | ----- | cu | 8G | In_Sync | FC Disabled |
| 400 | 5 | 16 | ----- | cu | 8G | In_Sync | FC Disabled |
| 401 | 5 | 17 | ----- | cu | 8G | In_Sync | FC Disabled |
| 402 | 5 | 18 | ----- | cu | 8G | In_Sync | FC Disabled |
| 403 | 5 | 19 | ----- | cu | 8G | In_Sync | FC Disabled |
| 404 | 5 | 20 | ----- | cu | 8G | In_Sync | FC Disabled |
| 405 | 5 | 21 | ----- | cu | 8G | In_Sync | FC Disabled |
| 406 | 5 | 22 | ----- | cu | 8G | In_Sync | FC Disabled |
| 407 | 5 | 23 | ----- | cu | 8G | In_Sync | FC Disabled |
| 408 | 5 | 24 | ----- | cu | 8G | In_Sync | FC Disabled |
| 409 | 5 | 25 | ----- | cu | 8G | In_Sync | FC Disabled |
| 410 | 5 | 26 | ----- | cu | 8G | In_Sync | FC Disabled |
| 411 | 5 | 27 | ----- | cu | 8G | In_Sync | FC Disabled |
| 412 | 5 | 28 | ----- | cu | 8G | In_Sync | FC Disabled |
| 413 | 5 | 29 | ----- | cu | 8G | In_Sync | FC Disabled |
| 414 | 5 | 30 | ----- | cu | 8G | In_Sync | FC Disabled |
| 415 | 5 | 31 | ----- | cu | 8G | In_Sync | FC Disabled |

SEE ALSO [bladeEnable](#), [chassisDisable](#), [chassisEnable](#), [portDisable](#), [portEnable](#), [slotShow](#), [switchEnable](#), [switchDisable](#), [switchShow](#)

bladeEnable

Enables all user ports on a blade.

Synopsis `bladeenable slot`

Description Use this command to enable all user ports on a blade. All ports within the blade that did not fail the power-on self-test (POST) are enabled (except for persistently disabled ports). They may come online if connected to a device, or remain offline if disconnected. Use the **bladeEnable** command to re-enable the blade after making configuration changes or running offline diagnostics.

If the switch is connected to a fabric through previously disabled ports, it rejoins the fabric. If this switch remains the principal switch at the end of the fabric countdown, it assigns itself a domain ID. If another switch assumes the principal role, the re-enabled switch becomes a subordinate switch and accepts a domain ID from the principal.

As each port is enabled, the front panel LED changes from a slow-flashing amber to nonflashing green for online ports or to nonflashing amber for ports that do not initialize. Disconnected ports remain unlit.

Notes You cannot disable a single blade when the entire chassis is disabled, or when the blade itself is faulted, powered off, or running diagnostics. Use **chassisEnable** to enable the ports on an entire chassis.

This command does not enable the ports for the given blade in a disabled logical switch.

Persistently disabled ports are not enabled by this command.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

slot
Specifies the slot number to be enabled.

Examples To display the slot status, enable the user ports in slot 4, and verify the settings:

```
switch:admin> slotshow
```

| Slot | Blade Type | ID | Status |
|------|------------|----|-------------------------------|
| 1 | SW BLADE | 2 | ENABLED |
| 2 | UNKNOWN | | VACANT |
| 3 | UNKNOWN | | VACANT |
| 4 | SW BLADE | 2 | ENABLED (User Ports Disabled) |
| 5 | CP BLADE | 1 | ENABLED |
| 6 | CP BLADE | 1 | ENABLED |
| 7 | SW BLADE | 2 | ENABLED |
| 8 | UNKNOWN | | VACANT |
| 9 | UNKNOWN | | VACANT |
| 10 | UNKNOWN | | VACANT |

```
switch:admin> bladeenable 4
Blade 4 is being enabled...Done
```

```
switch:admin> slotshow
```

| Slot | Blade Type | ID | Status |
|------|------------|----|---------|
| 1 | SW BLADE | 2 | ENABLED |
| 2 | UNKNOWN | | VACANT |
| 3 | UNKNOWN | | VACANT |
| 4 | SW BLADE | 2 | ENABLED |
| 5 | CP BLADE | 1 | ENABLED |
| 6 | CP BLADE | 1 | ENABLED |
| 7 | SW BLADE | 2 | ENABLED |
| 8 | UNKNOWN | | VACANT |
| 9 | UNKNOWN | | VACANT |
| 10 | UNKNOWN | | VACANT |

SEE ALSO [bladeDisable](#), [chassisDisable](#), [chassisEnable](#), [portEnable](#), [portDisable](#), [switchDisable](#), [switchShow](#)

bladeSwap

Swaps the area numbers for matching port pairs of two blades.

Synopsis `bladeSwap -src source_slot -dest destination_slot`

Description Use this command to swap the area numbers for matching port pairs of two blades. All ports must qualify for swapping for this command to succeed. It validates that the blades in the indicated slots are of the same type, have the same number of ports, and that the port pairs are in the same partition.

If all the ports qualify for swapping, this command automatically performs the following operations:

1. It enables the port swapping feature by executing the **portSwapEnable** command.
2. It takes all ports on both the source and destination blades offline by executing the **bladeDisable** command.
3. It swaps the matching port pairs on each of the specified blades.
4. It re-enables the blade by executing the **bladeEnable** command.

The result of this operation is persistent across reboots and power cycles.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command leaves both the source and destination blades in a disabled state. To enable all ports on a blade, issue the **bladeEnable** command. To enable individual ports, issue the **portEnable** command.

This command affects port swap information that is kept in its own database. It cannot be manipulated by editing the configuration database reported by **configShow**.

This command is not allowed if one or both the blades are faulty.

Operands This command has the following operands:

-src *source_slot* Specifies the slot number for the source blade.

-dest *destination_slot* Specifies the slot number for the destination blade. Use **slotShow** for a listing of valid slots.

Examples To swap area numbers between matching port pairs of two slots:

```
switch:admin> bladeswap -src 1 -dest 3
bladeswap done
```

SEE ALSO [portSwapEnable](#), [portSwapDisable](#), [portSwapShow](#), [portShow](#), [portEnable](#), [portDisable](#), [bladeEnable](#), [bladeDisable](#), [switchShow](#), [switchEnable](#)

bootLunCfg

Transparently configures the boot LUN for an HBA.

Synopsis

```
bootluncfg --add HBA_WWN PWWN LUN_ID
bootluncfg --delete HBA_WWN [PWWN LUN_ID]
bootluncfg --show
bootluncfg --help
```

Description Use this command to configure the boot LUN for an HBA.

Existing fabric-based boot LUN discovery allows the host's boot LUN information to be stored in the fabric zone database by using a zone name that contains the PWWN of an HBA port. The zone members consist of storage target PWWN and LUN ID.

This command provides a simplified and transparent procedure for configuring the boot LUN. Once configured, the HBA boot code queries the zone member list for the zone name matching the HBA PWWN to determine the boot target and LUN.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

HBA_WWN

Specifies the WWN of the HBA port connecting to the boot LUN. This is a hexadecimal number, for example, 10:00:00:000:05:1e:41:9a:cb.

PWWN

Specifies the WWN of the remote storage target's port. For example, 50:00:00:05:1e:41:9a:ca. This operand is required with the **--add** option and optional with the **--delete** option.

LUN_ID

Specifies the hexadecimal LUN identification. The LUN ID is represented as an eight-byte value (four-level LUN addressing), for example, 09AABBCCDDEEFF00. This operand is required with the **--add** option and optional with the **--delete** option.

--add

Configures the specified HBA (HBA_WWN) to boot from the specified LUN (LUN_WWN) accessible through the Target PWWN (PWWN).

--delete

Removes the mapping for the HBA (HBA_PWWN) to boot from the LUN (LUN_ID) accessible through the Target Port (PWWN).

--show

Displays all configured HBA to Port/LUN Mappings.

--help

Displays the command usage.

Examples

To configure a boot LUN:

```
switch:admin> bootluncfg --add 11:22:ab:44:44:ff:44:ca \  
1b:6c:55:55:55:3a:55:ff 9abc345fa1112410  
Operation Successful
```

To display existing Port/LUN mappings:

```
switch:admin> bootluncfg --show  
  
00:11:22:33:44:55:66:77  
00:00:00:00:aa:bb:cc:dd;00:00:00:01:ee:ff:11:22; \  
00:00:00:02:9a:bc:34:5f;00:00:00:03:a1:11:24:10  
  
aa:aa:aa:aa:aa:aa:aa:aa  
00:00:00:00:11:11:11:11;00:00:00:01:11:11:11:11; \  
00:00:00:02:9a:bc:34:5f;00:00:00:03:a1:11:24:10  
  
bb:aa:aa:aa:aa:aa:aa:aa  
00:00:00:00:11:11:11:11;00:00:00:01:11:11:11:11; \  
00:00:00:02:9a:bc:34:5f;00:00:00:03:a1:11:24:10
```

To remove an HBA to Port/LUN mapping:

```
switch:admin> bootluncfg --delete 11:22:ab:44:44:ff:44:ca \  
1b:6c:55:55:55:3a:55:ff 9abc345fa1112410  
Operation Successful
```

SEE ALSO

None

bufOpMode

Changes or displays the Buffer Optimized Mode.

Synopsis `bufopmode --set slot [-force | -f]`

`bufopmode --setall [-force | -f]`

`bufopmode --reset slot`

`bufopmode --resetall [-force | -f]`

`bufopmode --show slot`

`bufopmode --showall`

Description Use this command to display or change the buffer optimized mode on a switch.

When buffer optimized mode is enabled on a slot, additional buffers are allocated on the internal ports. Use this feature, if you have slow draining devices connected to the slot and there are no long distance links or F_Port buffers configured on that slot.

You must power off the slot before changing the buffer optimized mode. Changes take effect immediately after the slot is powered on.

Enabling buffer optimized mode removes all long distance configurations or F_Port buffers on that slot. You will be prompted if any F_Port buffers or long distance configurations are detected.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported on the Brocade FC16-32, FC16-48, and FC16-64 blades in a DCX 8510 chassis, and SX6 blade in a Brocade X6 Director. An asterisk (*) next to the slot configuration status indicates an unsupported blade. Refer to the example section for an illustration.

This command is not supported on Control Processor (CP) blades and core blades.

The **--set** operand is not supported in the following conditions on Brocade DCX 8510-family switches:

- On CP blade, core blade, or empty slot. This operand is allowed only when a blade is inserted into an empty slot and buffer optimized mode is enabled on all the other non empty slots.
- If the buffer optimized mode is already enabled on a slot in the chassis with 16 Gbps-capable core blades.
- If the buffer optimized mode is not enabled on any slot in the chassis with 16 Gbps-capable core blades. You can enable buffer optimized mode on all slots by using the **bufopmode --setall** command.

The **--reset** option is not supported on the control processor (CP) blades, core blades, and FC16 port blades.

The **--setall** and **--resetall** options are not supported on the FC8 core blades.

The 8G blades are supported in the 8G chassis. The 16G blades are supported on the 16G chassis. On 16G chassis, all blades, including the core blades, are enabled at the same time with the **--setall** option.

Operands This command has the following operands:

| | |
|---------------------------------|--|
| <i>slot</i> | Specifies the slot number. |
| --set [-force -f] | Enables buffer optimized mode on the specified slot. This command prompts for confirmation before removing any F_Port buffers or long distance configurations on that slot. Use the -force or -f option to execute this command without confirmation. |
| --setall [-force -f] | Enables buffer optimized mode on all slots. All switches on the chassis must be disabled using the chassisDisable command before executing this command. The switch must be rebooted for changes to take effect and therefore this operation is disruptive. The command prompts for confirmation before rebooting the switch. Use the -force or -f option to execute this command without confirmation. |
| --reset | Clears buffer optimized mode on the specified slot. |
| --resetall [-force -f] | Clears buffer optimized mode on all slots. The switch must be rebooted for changes to take effect and therefore this operation is disruptive. The command prompts for confirmation before rebooting the switch. Use the -force or -f option to execute this command without confirmation. |
| --show | Displays the current buffer optimized mode for the specified slot (On or Off). |
| --showall | Displays the current buffer optimized mode for all slots. An asterisk (*) next to the slot configuration status indicates an unsupported blade. |

Examples To display current buffer optimized mode for a single slot

```
switch:admin> bufopmode --show 12
Slot 12: buffer optimized mode - On
```

To enable buffer optimized mode for a given slot:

```
switch:admin> bufopmode --set 11
The slot must first be powered off to perform this operation

switch:admin> bufopmode --reset 12
The slot must first be powered off to perform this operation

switch:admin> slotpoweroff 11

switch:admin> slotpoweroff 12

switch:admin> bufopmode --set 12
Buffer optimized mode is turned on for slot 11

switch:admin> bufopmode --reset 12
```

Buffer optimized mode is turned off for slot 12

```
switch:admin> slotpoweron 11
```

```
switch:admin> slotpoweron 12
```

To display the changes:

```
switch:admin> bufopmode --showall
Slot 1: buffer optimized mode - Off
Slot 2: buffer optimized mode - Off
Slot 3: buffer optimized mode - Off
Slot 4: buffer optimized mode - On
Slot 9: buffer optimized mode - Off
Slot 10: buffer optimized mode - Off
Slot 11: buffer optimized mode - On
Slot 12: buffer optimized mode - Off
-----
* indicates buffer optimization not supported on this blade
```

To display the buffer optimized mode for all slots on a DCX 8510-8 with unsupported 16G port blades:

```
switch:admin> bufopmode --showall

Slot 1: buffer optimized mode - Off*
Slot 2: buffer optimized mode - Off*
Slot 3: buffer optimized mode - Off*
Slot 4: buffer optimized mode - Off*
Slot 5: buffer optimized mode - Off*
Slot 8: buffer optimized mode - Off*
Slot 9: buffer optimized mode - Off*
Slot 10: buffer optimized mode - Off*
Slot 11: buffer optimized mode - Off*
Slot 12: buffer optimized mode - Off*

-----
* indicates buffer optimization not supported on this blade
```

To enable buffer optimized mode on all slots:

```
switch:admin> chassisdisable

switch:admin> bufopmode --setall
Buffer optimized mode is turned on for slot 5
Buffer optimized mode is turned on for slot 8
Buffer optimized mode is turned on for slot 10
Buffer optimized mode is turned on for slot 12

WARNING: Must reboot switch for changes to take effect. Reboot now [y/n]? y
```

bufOpMode

To clear buffer optimized mode on all slots:

```
switch:admin> bufopmode --resetall
Buffer optimized mode is turned off for slot 5
Buffer optimized mode is turned off for slot 8
Buffer optimized mode is turned off for slot 10
Buffer optimized mode is turned off for slot 12
```

WARNING: Must reboot switch for changes to take effect. Reboot now [y/n]? y

SEE ALSO [slotShow](#), [slotPowerOn](#), [slotPowerOff](#)

cfgActvShow

Displays effective zone configuration information.

Synopsis `cfgactvshow`

Description Use this command to display the effective zone configuration information.

The current configuration is a single zone configuration that is currently in effect. The devices that an initiator sees are based on this configuration. The effective configuration is built when a specified zone configuration is enabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the effective zone configuration information:

```
switch:admin> cfgactvshow
Effective configuration:
cfg:    c4
zone:  z3      33:07:06:05:04:03:02:01
zone:  z4      44:01:23:45:67:89:a0:bc
                40:01:23:45:67:89:a0:bc
```

SEE ALSO [cfgClear](#), [cfgDelete](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#)

cfgAdd

Adds a member to a zone configuration.

Synopsis `cfgadd "cfgName", "member[; member...]"`

Description Use this command to add one or more members to an existing zone configuration.

This command changes the Defined Configuration. For the change to take effect, enable the configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"cfgName"

Specify a name for the zone configuration, enclosed in double quotation marks.

"member"

Specify a zone member or a list of zone members to be added to the configuration. The list must be enclosed in double quotation marks. Members must be separated by semicolons.

Examples To add two new zones to the configuration "Test_cfg":

```
switch:admin> cfgadd "Test_cfg", "greenzone;bluezone"
```

SEE ALSO [cfgClear](#), [cfgCreate](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#)

cfgClear

Clears all zone configurations.

Synopsis `cfgclear`

Description Use this command to clear all zone information in the transaction buffer. All defined zone objects in the transaction buffer are deleted. If an attempt is made to commit the empty transaction buffer while a zone configuration is enabled, you are warned to first disable the enabled zone configuration or to provide a valid configuration with the same name.

After clearing the transaction buffer with the **cfgClear** command, use the **cfgDisable** command to clear the entire zoning configuration (both the transaction buffer and the active configuration). If there is no current active zoning configuration, or you just want to clear the transaction buffer, use the **cfgSave** command to commit the transaction.

If the default zone access mode is "No Access", this command recreates the default zoning objects.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands None

Examples To clear all zones and then clear nonvolatile memory:

```
switch:admin> cfgclear
The Clear All action will clear all Aliases, Zones, FA Zones
and configurations in the Defined configuration.
Run cfgSave to commit the transaction or cfgTransAbort to
cancel the transaction.
Do you really want to clear all configurations? \
  (yes, y, no, n): [no] n
```

```
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled. Until the Effective
configuration is re-enabled, merging new switches into the
fabric is not recommended and may cause unpredictable
results with the potential of mismatched Effective Zoning
configurations.
Do you want to save Defined zoning configuration only? \
  (yes, y, no, n): [no]n
```

SEE ALSO [cfgDisable](#), [cfgEnable](#), [cfgSave](#)

cfgCreate

Creates a zone configuration.

Synopsis `cfgcreate "cfgName", "member[;member...]"`

Description Use this command to create a new zone configuration.

This command changes the Defined Configuration (see **cfgShow**). For the change to become effective, enable the configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the **cfgSave** command.

Refer to the **zoneCreate** command for more information on *name* and *member* specifications.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"cfgName"

Specify a name for the zone configuration in double quotation marks. A zone configuration name must begin with a letter followed by any number of letters, numbers, and underscores. Names are case-sensitive. For example, "Cfg_1" and "cfg_1" are different zone configurations. Blank spaces are ignored.

"member"

Specify a zone member or list of zone members to be added to the configuration. The list must be enclosed in double quotation marks. Members must be separated by semicolons. The zone configuration member list must have at least one member. Empty member lists are not allowed.

Examples To create a configuration containing three zones:

```
switch:admin> cfgcreate "USA_cfg", \  
"Purple_zone;Blue_zone;Green_zone"
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#)

cfgDelete

Deletes a zone configuration.

Synopsis `cfgdelete "cfgName"`

Description Use this command to delete a zone configuration.

This command changes the Defined Configuration (see **cfgShow**). For the change to become effective, enable the configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operand is required:

"cfgName"

Specify a name for the zone configuration to be deleted. The name must be enclosed in double quotation marks.

Examples To delete a zone configuration:

```
switch:admin> cfgdelete "USA_cfg"
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgCreate](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#)

cfgDisable

Disables a zone configuration.

Synopsis `cfgdisable`

Description Use this command to disable the current zone configuration. The fabric returns to non-zoning mode, a mode in which all devices can see each other.

This command ends and commits the current zoning transaction buffer to both volatile and nonvolatile memory. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message displays on the other switches to indicate that the transaction was aborted.

If the default zone access mode is "No Access", then this command becomes equivalent to **cfgEnable** "d_efault_Cfg". Refer to **defZone** help for information on zone access configuration.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands None

Examples To disable the current zone configuration:

```
switch:admin> cfgdisable
You are about to disable zoning configuration. This
action will disable any previous zoning configuration.
Do you want to disable zoning configuration? \
  (yes, y, no, n): [no] y
```

SEE ALSO [cfgClear](#), [cfgEnable](#), [cfgSave](#)

cfgEnable

Enables a zone configuration.

Synopsis `cfgenable "cfgName"`

Description Use this command to enable a zone configuration. The command builds the specified zone configuration. It checks for undefined zone names, zone alias names, or other inconsistencies, by expanding zone aliases, removing duplicate entries, and then installing the effective configuration.

If the build fails, the previous state is preserved (zoning remains disabled, or the previous effective configuration remains in effect). If the build succeeds, the new configuration replaces the previous configuration. Refer to the **cfgShow** command for a description of the defined and effective configurations.

If there are open transactions in the fabric, only single transaction can be saved. Use the **cfgtransshow --opentrans** command to view the list of all the domains in the fabric with open transactions.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operand is required:

"cfgName"

Specifies the name of the zone configuration. The name must be enclosed in double quotation marks.

Examples To enable the zone configuration USA_cfg:

```
switch:admin> cfgenable "USA_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'USA_cfg' configuration \
    (yes, y, no, n): [no] y
zone config "USA_cfg" is in effect
Updating flash ...
```

To enable the zone configuration USA_cfg when there are open transactions in the fabric:

```
switch:admin> cfgenable "USA_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Multiple open transactions are pending in this fabric. Only one
transaction can be saved. Please abort all unwanted transactions
using the cfgtransabort command. Use the cfgtransshow --opentrans
```

cfgEnable

```
command to display a list of domains with open transactions
switch:admin> cfgtransabort 271010736
switch:admin> cfgenable "USA_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'USA_cfg' configuration \
    (yes, y, no, n): [no] y
zone config "USA_cfg" is in effect
Updating flash ...
```

SEE ALSO [cfgClear](#), [cfgDisable](#), [cfgSave](#), [cfgShow](#)

cfgRemove

Removes a member from a zone configuration.

Synopsis `cfgremove "cfgName","member[; member...]"`

Description Use this command to remove one or more members from an existing zone configuration. If all members are removed, the zone configuration is deleted.

This command changes the Defined Configuration (see **cfgShow**). For the change to become effective, enable the configuration with the **cfgEnable** command. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the **cfgSave** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"cfgName"

Specify a name for the zone configuration, enclosed in double quotation marks.

"member"

Specify a zone member or a list of zone members to be removed from the configuration. The list must be enclosed in double quotation marks. Members must be separated by semicolons.

Examples To remove a zone from a configuration:

```
switch:admin> cfgremove "Test_cfg", "bluezone"
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgCreate](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgSave](#), [cfgShow](#), [cfgTransAbort](#), [cfgTransShow](#)

cfgSave

Saves the zone configuration to nonvolatile memory.

Synopsis `cfgsave`

Description Use this command to save the current zone configuration. This command writes the defined configuration and the name of the effective configuration to nonvolatile memory in all switches in the fabric.

The saved configuration is automatically reloaded at power on, and, if a configuration was in effect at the time it was saved, the same configuration is reinstalled with an automatic **cfgEnable** command.

Because the saved configuration is reloaded at power on, only valid configurations are saved. The **cfgSave** command validates the effective configuration by performing the same tests as the **cfgEnable** command. If the tests fail, an error message is displayed and the configuration is not saved.

This command ends and commits the current transaction. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted. A message displays on the other switches to indicate that the transaction was aborted.

If there are open transactions in the fabric, only single transaction can be saved. Use the **cfgtransshow --opentrans** command to view the list of all the domains in the fabric with open transactions.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands None

Examples To save a zone configuration:

```
switch:admin> cfgsave
WARNING!!!
The changes you are attempting to save will render the
Effective configuration and the Defined configuration
inconsistent. The inconsistency will result in different
Effective Zoning configurations for switches in the fabric if
a zone merge or HA failover happens. To avoid inconsistency
it is recommended to commit the configurations using the
cfgenable command.
Do you want to save Defined zoning configuration
only? (yes, y, no, n): [no] y
```

To save a zone configuration if there are multiple open transactions:

```
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Multiple open transactions are pending in this fabric. Only one
transaction can be saved. Please abort all unwanted transactions
using the cfgtransabort command. Use the cfgtransshow --opentrans
command to display a list of domains with open transactions
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no] n
```



```
switch:admin> cfgtransabort 271010736
switch:admin> cfigsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
If the update includes changes to one or more traffic isolation
zones, you must issue the 'cfgenable' command for the changes
to take effect.
Do you want to save the Defined zoning configuration only? (yes, y, no, n): [no] y
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgCreate](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgShow](#), [cfgTransAbort](#), [cfgTransShow](#)

cfgShow

Displays zone configuration information.

Synopsis `cfgshow [--ic] ["pattern"] [, mode]`

`cfgshow [--transdiffs]`

`cfgshow [--transdiffsonly]`

`cfgshow --help`

Description Use this command to display zone configuration information.

If no operand is specified, all zone configuration information (both defined and effective) is displayed. If the local switch has an outstanding transaction, this command displays the most recently edited zone configuration that has not yet been saved. If the local switch has no outstanding transaction, this command displays the committed zone configuration.

If a pattern is specified, only matching configurations are displayed.

- The **defined configuration** is the complete set of all zone objects that have been defined in the fabric. There can be multiple zone configurations defined, but only one can be enabled at a time. There might be inconsistencies in the definitions, zones, or aliases that are referenced but not defined, or there might be duplicate members. The defined configuration is the current state of the administrator input.
- The **effective configuration** is the single zone configuration that is currently enabled. The devices that an initiator sees in the fabric are based on this configuration. The effective configuration is built when a specific zone configuration is enabled and all error checking has been completed successfully.

Use the `--transdiffs` and `--transdiffsonly` options to view changes in the current transaction.

The command output displays changes in the current transaction by the following notations:

- An asterisk(*) before any tag indicates a change in that zone, zone configuration, alias or any other entity in the zone configuration.
- A plus(+) before any entity indicates that it is a newly added entity.
- A minus(-) before any entity indicates that it is a deleted entity.

When this command is executed after a zoning transaction was aborted on the local switch, it displays the following warning message: "Warning: Current Zoning Transaction was aborted. Reason code = Zone Config update received."

When default zoning is enabled with "No Access" mode, "No Effective configuration: (No Access)" is displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operands are optional:

`--ic`

Displays all zone configuration names for a given pattern without case distinction.

"pattern"

A regular expression that matches zone configuration names. The pattern must be enclosed in quotation marks and can contain the following:

Question mark (?)

Matches any single character

Asterisk (*)

Matches any string of characters.

Range

Matches any character within the range. Ranges must be enclosed in square brackets, for example, [0-9] or [a-f].

mode

Specify **0** to display the contents of the transaction buffer (the contents of the current transaction), **1** to display the contents of committed defined database, and **2** to display the contents of effective zone configuration. The default value is 0.

--transdiffs

Displays changes in the current transaction.

--transdiffsonly

Displays only the changes in the current transaction.

Examples

To display all zone configurations that start with "Test":

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

To display all zone configurations that start with "Test", regardless of the case:

```
switch:admin> cfgshow -ic "Test*"
cfg:   Test1 Blue_zone
cfg:   TEST2_cfg Red_zone; Blue_zone
```

To display all zone configuration information:

```
switch:admin> cfgshow
Defined configuration:
cfg:   USA1   Blue_zone
cfg:   USA_cfg Red_zone; Blue_zone
zone:  Blue_zone
      1,1; array1; 1,2; array2
zone:  Red_zone
      1,0; loop1
alias: array1  21:00:00:20:37:0c:76:8c; \
      21:00:00:20:37:0c:71:02
alias: array2  21:00:00:20:37:0c:76:22; \
      21:00:00:20:37:0c:76:28
alias: loop1   21:00:00:20:37:0c:76:85; \
      21:00:00:20:37:0c:71:df

Effective configuration:
cfg:   USA_cfg
zone:  Blue_zone
      1,1
```

cfgShow

```
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
zone: Red_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df
```

To display only configuration names:

```
switch:admin> cfgshow ""
cfg: USA1 Blue_zone
cfg: USA_cfg Red_zone; Blue_zone
```

To display changes in the current transaction:

```
switch:admin> cfgshow --transdiffs
Defined configuration:
cfg: fabric_cfg Blue_zone

zone: Blue_zone
1,1; array1; 1,2; array2

*zone: green_zone
-1,1; 1,2; +6, 15

*zone: +red_zone
5,1; 4,2

alias: array1 21:00:00:20:37:0c:76:8c; \
21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; \
21:00:00:20:37:0c:76:28

Effective configuration:
cfg: fabric_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
```

To display only the changes in the current transaction:

```
switch:admin> cfgshow --transdiffsonly
*zone: -Blue_zone
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
```

```
*zone: green_zone
      1,1; -1,2;+5,4;+ 21:00:00:20:37:0c:76:55

*zone: +red_zone
      5,4; 5,6
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgCreate](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgTransAbort](#), [cfgTransShow](#)

cfgSize

Displays zone database size details.

Synopsis `cfgsize [integer]`

Description Use this command to display the size details of the zone database.

The size details include maximum size, the committed size, and the transaction size of the Zone database in bytes.

Zone DB maximum size

Defines the upper limit for the Zone defined configuration, determined by the amount of nonvolatile memory available for storing the defined configuration. The Zone DB maximum size is further reduced due to a message header that is propagated with the zone configuration to all switches in the fabric.

The maximum supported zone database size is 2 MB for chassis based switches and 1 MB for other switches.

Available Zone DB size

Displays the size of the database that is actually available for storage.

Committed size

Displays the size of the defined configuration currently stored in nonvolatile memory.

Transaction size

Displays the size of the uncommitted defined configuration. This value will be nonzero if the defined configuration is being modified by Telnet, API, and so forth; otherwise it is 0.

Refer to the **cfgShow** help page for a description of defined and effective zone configurations.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

In Virtual Fabric environments, the zone database size of all partitions taken together cannot exceed 2 MB for chassis based switches and 1 MB for other switches. Exceeding the zone database size limit can result in unexpected behavior. Note that in Virtual Fabric mode, the zone database limit is enforced per partition, but not chassis-wide.

Operands The following operand is optional:

integer

If a nonzero integer is specified, the size of the nonvolatile memory allocated for the zone database is displayed. The zone database includes both the defined and effective configurations. The database size is displayed in bytes.

Examples To display the zone database on a Brocade 6510

```
switch:admin> cfgsize

Zone DB max size - 1045274 bytes
Available Zone DB size - 1041503 bytes

    committed - 2759
    transaction - 0
```

SEE ALSO [ad.cfgShow](#), [zoneHelp](#)

cfgTransAbort

Aborts the current zoning transaction.

Synopsis `cfgtransabort [token]`

Description Use this command to abort the current zoning transaction without committing it. All changes made since the transaction was started are removed and the zone configuration database is restored to the state before the transaction was started.

If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch remains open and unaffected.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands When invoked without an operand, this command aborts the current transaction. The following operand is optional:

token

Specify the token ID of the transaction to be aborted. Use the **cfgTransShow** command to obtain the token ID of a transaction.

Examples To abort the current transaction:

```
switch:admin> cfgtransabort
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgCreate](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#), [cfgTransShow](#)

cfgTransShow

Displays information about the current zoning transaction.

Synopsis `cfgtransshow`

`cfgtransshow --opentrans`

`cfgtransshow --help`

Description Use this command to display the ID of the current zoning transaction. In addition, the command provides information on whether or not the transaction can be aborted. The transaction cannot be aborted if it is an internal zoning transaction.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands This command has the following operands:

--opentrans

Displays local open transaction token details and the list of domains with open transactions.

--help

Displays the command usage.

Examples To display the current transaction:

```
switch:admin> cfgtransshow
There is no outstanding zone transaction
```

```
switch:admin> cfgclear
Do you really want to clear all configurations? \
  (yes, y, no, n): [no] y
Clearing All zoning configurations...
```

```
switch:admin> cfgtransshow
Current transaction token is 271010736
It is abortable
```

cfgTransShow

To display the current transaction and the list of domains with open transactions:

```
switch:admin> cfgtransshow --opentrans  
Current transaction token is 0x3109  
It is abortable  
  
Transactions Detect: Capable  
  
Current Open Transactions  
Domain List:  
-----  
1 2 3 4
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgCreate](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#), [cfgTransAbort](#)

chassisBeacon

Sets chassis beaconing mode.

Synopsis `chassisbeacon [mode]`

Description Use this command to enable or disable beaconing on a chassis. Chassis beaconing can be used to locate a failing chassis in a group of chassis. Use the **portBeacon** command to locate a failing port, and use the **switchBeacon** command to locate a failing (logical) switch.

When beaconing mode is turned on, the port LEDs flash green at various rates across the chassis. The beaconing continues until you turn it off.

Beaconing mode takes over the port LEDs. The normal flashing LED pattern associated with an active, faulty, or disabled port is suppressed, and only the beaconing pattern is shown. Other commands are still executable and functional. However, if diagnostic frame-based tests such as **portLoopbackTest** are executed, the diagnostic LED pattern is interleaved with the beaconing pattern.

The **chassisBeacon** command is one of the commands that control beaconing. Each command has a clearly defined scope of action:

- The **portBeacon** command enables or disables beaconing on a specified port.
- The **switchBeacon** command enables or disables beaconing on all ports in the current logical switch.
- The **chassisBeacon** command enables or disables beaconing on all ports in the chassis.
- The **portPeerBeacon** command enables or disables beaconing to identify the interconnections between ports.

The actions of the beaconing commands are independent and mutually exclusive. For example, if you enabled beaconing on the logical switch and you want to enable beaconing on the entire chassis, you must first disable switch beaconing with the **switchBeacon** command before you can use the **chassisBeacon** command to enable beaconing on the entire chassis. Likewise, existing **portBeacon** settings remain unaffected if you enable or disable beaconing on the switch or on the chassis. Failure to disable existing beaconing commands before using a different type of beaconing may cause the commands to interfere with each other in unexpected ways.

On a Brocade X6 Director, enabling chassis beaconing activates beaconing on CP blades also.

To determine whether or not beaconing is enabled or disabled on the switch or chassis, use the **switchBeacon** or **chassisBeacon** command without operands. A value of 0 indicates that the command is disabled, a value of 1 indicates that the command is enabled. Issue the **portBeacon --show** command to display beaconing for a specific port. The **switchShow** command displays the status of the **switchBeacon** command only.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

mode

Specify 1 to enable chassis beaconing mode or 0 to disable beaconing mode. This operand is optional. If no operand is specified, the current value is displayed.

chassisBeacon

Examples To turn chassis beaconing mode on:

```
switch:admin> chassisbeacon 1
Chassis beacon success 1
```

To turn beaconing mode off:

```
switch:admin> chassisbeacon 0
Chassis beacon success 0
```

To display the chassis beaconing mode:

```
switch:admin> chassisbeacon
Value = 0
```

SEE ALSO [portBeacon](#), [switchShow](#), [switchBeacon](#), [portPeerBeacon](#)

chassisCfgPerrThreshold

Configures parity error threshold and fault capability on a blade or switch.

Synopsis `chassiscfgperrthreshold --enable [-threshold threshold_value]`

`chassiscfgperrthreshold --disable`

`chassiscfgperrthreshold --show`

`chassiscfgperrthreshold --help`

Description Use this command to enable or disable the fault capability on a blade and switch based on the parity error threshold, to set the threshold value for parity errors, or to display the configuration. If the fault capability is enabled, the state of the blade or switch changes to Faulty when the parity errors exceed the threshold value within a time interval of 5 minutes. The default threshold value is 2. You can modify the threshold using the **-threshold *threshold_value*** option.

The configuration is persistent across system reboots and high availability (HA) failover.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--enable

Enables the fault capability on a blade or switch. The default threshold value is 2. The following operand is optional:

-threshold *threshold_value*

Specifies the threshold value for parity errors. The range is from 2 through 100.

--disable

Disables the fault capability on a blade or switch.

--show

Displays the configured threshold value and the current status of the fault capability as enabled or disabled.

--help

Displays the command usage.

Examples To enable fault capability:

```
switch:admin> chassiscfgperrthreshold --enable
```

To disable fault capability:

```
switch:admin> chassiscfgperrthreshold --disable
```

To enable fault capability and set the threshold value to 76:

```
switch:admin> chassiscfgperrthreshold --enable -threshold 76
```

chassisCfgPerrThreshold

To display the configuration:

```
switch:admin> chassiscfgperrthreshold --show  
Configured Threshold is 76 and the Fault is enabled.
```

SEE ALSO None

chassisDisable

Disables all user ports in a chassis.

Synopsis `chassisdisable [-force]`

Description Use this command to disable a Virtual Fabric-aware chassis. All Fibre Channel ports are taken offline. This command prompts for confirmation unless the **-force** option is used. If the chassis is partitioned into logical switches that are part of logical fabrics, the remaining switches in these fabrics reconfigure. As each port is disabled, the front panel LED changes to a slow-flashing amber.

You must disable the chassis before making configuration changes or running offline diagnostic tests. Commands that require the chassis to be disabled generate an error message if invoked while the chassis is enabled. It is not necessary to disable a chassis before rebooting or powering off the switch.

To disable the ports of a single logical switch, use the **switchDisable** command. To disable the ports of a single blade, use the **bladeDisable** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands **-force** Disables the chassis without confirmation. This operand is optional.

Examples To disable a chassis:

```
switch:admin> chassisdisable
This command can cause disruption to\
multiple logical switches.
Are you sure you want to disable all chassis ports now?\
(yes, y, no, n): [no]y
```

SEE ALSO [bladeDisable](#), [bladeEnable](#), [chassisEnable](#), [switchShow](#), [switchDisable](#), [switchEnable](#), [switchCfgPersistentEnable](#), [switchCfgPersistentDisable](#)

chassisDistribute

Distributes IP filter policies.

Synopsis `chassisdistribute -db ipfilter -fid target_FID [-force]`

`chassisdistribute -db ipfilter -domain target_domain_list [-force]`

Description

Use this command to manually distribute the IP Filter policy database. The behavior of this command depends on whether Virtual Fabrics are enabled or disabled.

- If Virtual Fabrics are enabled on the originating switch, this command distributes the IP filter policy database to all chassis that are connected through the specified logical switch (FID). If the target chassis has Virtual Fabrics enabled, the IP filter policy database is distributed to all logical switches configured on the target. If Virtual Fabrics are not enabled on the target, the IP filter policy database is distributed to the default logical switch of the target.
- If Virtual fabrics are disabled on the originating switch, the IP filter policy database is distributed to a list of target switches specified by their Domain IDs. If a target domain has Virtual Fabrics enabled, the IP filter policy database is distributed to all logical switches on the target domain. If Virtual Fabrics are not enabled on the target domain, the IP filter policy database is distributed to the default logical switch of the target domain.

This command distributes the entire IP Filter database and overwrites any existing IP Filter policies on the receiving switches. The **chassisDistribute** command does not enforce FCS policy.

The target chassis or switches must be capable of accepting the distribution. The distribution is aborted if one of the connected chassis or domains is configured to reject the distribution. Use the **fddCfg** command to configure the fabric-wide policies that control distribution behavior.

When upgrading from pre-7.2.0 to 7.2.0 or later versions, the RPC port rules will be removed from the IP Filter policy database.

The command output includes the following information:

FID

Indicates whether Virtual Fabrics are enabled or not on the domain that receives the distribution. If Virtual Fabrics are enabled, the FID is displayed, and all logical switches which are part of the chassis are targeted to receive the distribution.

DOMAIN

Displays the Domain ID of the originating switch.

CHASSISWWN

Displays the WWN of the originating or receiving chassis. Displays the originating switch WWN in switches running Fabric OS versions that do not support the chassis WWN feature.

CHASSISNAME

Displays the name of the chassis.

SUPPORTED_DIST

Displays yes if the distribution is supported on the receiving chassis. Displays no if the distribution is not supported.

Notes This command distributes the IP Filter database only. To distribute other security databases, use the **distribute** command.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-db ipfilter

Specifies the policy database to be distributed with this command. The only accepted value for **-db** is **ipfilter**. The database name is case insensitive. The legacy notation ("IPFILTER") still works but is not necessary. The following operands are mutually exclusive:

-fid target_FID

Specifies the target logical switch for the distribution. The logical switch is identified by its fabric ID (FID). When a FID is specified, the database is distributed to all chassis that are connected to the specified logical switch (FID) and are configured to accept the distribution. This operand is required and valid only when the originating switch is in Virtual Fabric mode.

-domain target_domain_list

Specifies a list of target switch domain IDs that should receive the database. The list must be enclosed in double quotation marks; members must be separated by a semicolon. A wildcard (*) may be specified to include all switches in the fabric that support the distribute feature. This operand is required and valid only when Virtual Fabrics are disabled.

-force

Distributes the database without confirmation.

Examples To distribute the IP Filter policies to all chassis that are connected through the logical fabric 128 and support the distribute feature:

```
switch:admin> chassisDistribute -db ipfilter -fid 128
FID   DOMAIN   CHASSISWWN           CHASSISNAME  SUPPORTED_DIST
-----
 128   3  10:00:00:05:1e:38:ac:0e  DCX_93       yes
NON-VF 4  10:00:00:05:1e:39:bd:0f  Brocade300   yes
 128   98 10:00:00:05:1e:41:22:9f  Brocade6510  no
```

```
chassisDistribute will distribute the database(s) to \
above topology.
Would you like to continue [Y/N] : y
ChassisDistribute operation succeeded for above topology
```

To distribute the IP Filter policies to all switches in the fabric that support the distribute feature (Virtual Fabrics are disabled on the evoking switch):

```
switch:admin> chassisDistribute -db ipfilter -domain ""
FID   DOMAIN   CHASSISWWN           CHASSISNAME  SUPPORTED_DIST
-----
 128   3  10:00:00:05:1e:38:ac:0e  DCX_93       no
NON-VF 4  10:00:00:05:1e:39:bd:0f  Brocade300   yes
```

chassisDistribute

```
128  98 10:00:00:05:1e:41:22:9f Brocade6510      no
ChassisDistribute will distribute the database(s) to \
above topology.
Would you like to continue [Y/N] : y
ChassisDistribute operation succeeded for above topology
```

SEE ALSO [distribute](#), [fddCfg](#)

chassisEnable

Enables all user ports in a chassis.

Synopsis `chassisenable`

Description Use this command to enable a Virtual Fabric-aware chassis. All Fibre Channel ports that passed the power-on self test (POST) are enabled. They may come online if connected to a device, or remain offline if disconnected. Use **chassisEnable** to re-enable the chassis after making configuration changes or running offline diagnostics.

If the chassis is partitioned into multiple logical switches and physically connected to multiple logical fabrics, the logical switches rejoin their fabrics.

As each port is enabled, the front panel LED changes from slow-flashing amber to nonflashing green for online ports, or to nonflashing amber for ports that do not initialize. Disconnected ports remain unlit. Loopback ports slowly flash green when online.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable a chassis:

```
switch:admin> chassisenable
```

SEE ALSO [bladeDisable](#), [bladeEnable](#), [chassisDisable](#), [switchShow](#), [switchDisable](#), [switchEnable](#), [switchCfgPersistentEnable](#), [switchCfgPersistentDisable](#)

chassisName

chassisName

Displays or sets the chassis name.

Synopsis `chassisname [name]`

Description Use this command to display or change the name associated with the chassis. The configured chassis name appears in all the chassis event RASLog messages.

Use this command without parameters to display the current chassis name. Use this command with the *name* operand to assign a new chassis name.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

name

Specify a new name for the chassis, optionally in double quotation marks. A chassis name can include up to 31 characters. A chassis name must begin with a letter, and can consist of letters, numbers, underscore or hyphen characters. Spaces are not permitted.

Examples To change the chassis name to "dilbert":

```
switch:admin> chassisname dilbert
switch:admin> chassisname
dilbert
```

SEE ALSO [switchName](#)

chassisShow

Displays all field replaceable units (FRUs).

Synopsis `chassisshow`

Description Use this command to display the Field Replaceable Unit (FR) header content for each object in the chassis and chassis backplane version. This command displays the following information:

1. The chassis family, for example, DCX 8510.
2. The backplane version number, in hexadecimal: Chassis Backplane Revision: *xx*
3. The first line of each record contains the object ID. If the FRU is part of an assembly, a brief description is displayed in parentheses.
 - a. Object type: CHASSIS, FAN, POWER SUPPLY, SW BLADE (switch), CP BLADE (control processor), WWN (world wide name), or UNKNOWN.
 - b. Object number: Slot *nn* (for blades), Unit *nn* (for everything else).
4. FRU header version number: Header Version: *x*
5. Displays the maximum allowed power consumption for a given hardware component: positive for power supplies and negative for power consumers. The combined total maximum allowed power consumption for the entire chassis is shown in the output of the `slotShow -p` command. Power Consume Factor: *-xxx*
6. Displays the real-time power consumption for each FRU that supports real-time power measurement. Only 16Gb blades currently support real time measurement of power being consumed by these blades. When the capability is not supported for a FRU, the line is suppressed. Power Usage (Watts):
7. Factory part number (up to 14 characters): Factory Part Num: *xx-xxxxxx-xx*
8. Factory serial number (up to 12 characters): Factory Serial Num: *xxxxxxxxxx*
9. FRU manufacture date: Manufacture: Day: *dd* Month: *mm* Year: *yyyy*
10. Date of the last FRU header update: Update: Day: *dd* Month: *mm* Year: *yyyy*
11. Cumulative number of days the FRU has been powered on: Time Alive: *dddd* days
12. Time elapsed, in days, since the FRU was last powered on: Time Awake: *dddd* days
13. Externally supplied ID (up to 10 characters): ID: *xxxxxxxxxx*
14. Externally supplied part number (up to 20 characters): Part Num: *xxxxxxxxxxxxxxxxxxxxxx*
15. Externally supplied serial number (up to 20 characters): Serial Num: *xxxxxxxxxxxxxxxxxxxxxx*
16. Externally supplied revision number (up to 4 characters): Revision Num: *xxxx*

The output of this command depends on the platforms on which it is executed. On some platforms, for certain FRU types, a few items may not be available. In these cases, the lines are suppressed. Possibly affected lines are 1, 3 through 7, 9, and 11 through 14. In addition, for lines 11 through 14, if there is no data set, these lines are suppressed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the FRUs for a DCX 8510-8:

```
switch:user> chassisshow

Chassis Family:          DCX8510-8
Chassis Backplane Revision: 0

SW BLADE Slot: 1
Header Version:          2
Power Consume Factor:    -180
Power Usage (Watts):     -286
Factory Part Num:        60-1002144-02
Factory Serial Num:      BQB0345F00G
Manufacture:             Day: 9  Month: 11  Year: 2010
Update:                  Day: 19  Month: 2  Year: 2011
Time Alive:              41 days
Time Awake:              1 days

SW BLADE Slot: 2
Header Version:          2
Power Consume Factor:    -180
Power Usage (Watts):     -306
Factory Part Num:        60-1002144-02
Factory Serial Num:      BQB0345F02R
Manufacture:             Day: 21  Month: 11  Year: 2010
Update:                  Day: 19  Month: 2  Year: 2011
Time Alive:              41 days
Time Awake:              1 days

SW BLADE Slot: 3
Header Version:          2
Power Consume Factor:    -180
Power Usage (Watts):     -315
Factory Part Num:        60-1002144-02
Factory Serial Num:      BQB0345F01N
Manufacture:             Day: 16  Month: 11  Year: 2010
Update:                  Day: 19  Month: 2  Year: 2011
Time Alive:              39 days
Time Awake:              1 days
(output truncated)
```

SEE ALSO [slotShow](#)

classConfig

Displays RBAC class permissions.

Synopsis `classconfig --show class_name | -all | -classlist`

`classconfig --showcli command`

`classconfig --showroles class_name`

`classconfig --help`

Description

Use this command to display information about role-based access control (RBAC) permissions for one or all meta-object format (MOF) classes, to display permissions for a specified command, or to display the permissions for a specified MOF class.

Fabric OS commands are grouped into feature sets called MOF classes. For example, the commands **ldapCfg**, **passwd**, **passwdCfg**, and **userConfig** are all related to User Management and are therefore grouped together under a MOF class called **UserManagement**.

The pre-defined roles Root, Factory, Admin, User, SwitchAdmin, ZoneAdmin, FabricAdmin, BasicSwitchAdmin, SecurityAdmin and Operator provide a mechanism for further restricting access to commands of a certain class by role-specific permissions. For example, a user with the ZoneAdmin role will have access to the commands under the MOF class Zoning but not to those under the UserManagement class.

The following RBAC permissions are supported in Fabric OS:

- O = observe
- OM = observe-modify
- N = no access

Note that the MOF class level permissions extend to all commands in that class, but not necessarily to each and every command option. For example, a command may have the RBAC class permission of "OM", but a show only option under that command may have the permission "O". At the role level, a certain role may be excluded from viewing the command information, in which case the permission for that role would be "N".

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following optional operand:

--show

Displays RBAC permission information. One of the following operands is required:

***class_name* or -all**

Displays permissions for a single MOF class or for all classes. The output displays the class name, the CLIs grouped under the specified class, the command options, and the RBAC class permissions for each command option.

-classlist

Displays an alphabetical listing of all MOF classes supported in Fabric OS.

--showcli *command*

Displays the RBAC permissions for the specified command and associated command options. The output displays the command name, the command options, the RBAC class permissions for each command option, the MOF class, and the context restriction for the command in VF mode only. If the length of a row is too long to fit in one line, the strings displayed under RBAC Class and Option column will be split and displayed in two lines.

--showroles *class_name*

Displays the role permissions the specified MOF class.

--help

Displays the command usage.

Examples

To display an alphabetical listing of all MOF classes supported in Fabric OS:

```
switch:admin> classconfig --show -classlist
List of RBAC classes:
 1. AdminDomains
 2. ADSelect
 3. AG
 4. APM
 5. Audit
 6. Authentication
 7. Blade
 8. ChassisConfiguration
 9. ChassisManagement
10. ConfigManagement
11. Configure
12. DCE
13. Debug
14. Diagnostics
15. DMM
16. EncryptionConfiguration
17. EncryptionManagement
18. EthernetConfig
19. Fabric
20. FabricDistribution
21. FabricRouting
22. FabricWatch
23. Factory
24. FCoE
25. FICON
26. FIPSBootprom
27. FIPSCfg
28. FirmwareKeyManagement
29. FirmwareManagement
30. FRUManagement
31. HA
32. IPfilter
33. IPSec
34. ISCSI
35. License
36. LocalUserEnvironment
37. Logging
38. LogSupportsave
39. ManagementAccessConfiguration
40. ManagementServer
41. NameServer
```


42. Nocheck
43. NxPortManagement
44. PhysicalComputerSystem
45. PKI
46. PortMirror
47. RADIUS
48. Reboot
49. Restricted
50. RoleConfig
51. RoutingAdvanced
52. RoutingBasic
53. Security
54. SessionManagement
55. SNMP
56. Statistics
57. StatisticsDevice
58. StatisticsPort
59. SwitchConfiguration
60. SwitchManagement
61. SwitchManagementIPConfiguration
62. SwitchPortConfiguration
63. SwitchPortManagement
64. Topology
65. USBManagement
66. UserManagement
67. WWNCard
68. Zoning

To display the RBAC permissions for the commands included in the UserManagement class:

```
switch:admin> classconfig --show UserManagement
```

```
RBAC class Name : UserManagement
```

| CLI | Option | Permission |
|------------|---------------------|------------|
| ldapcfg | help | O |
| ldapcfg | maprole | OM |
| ldapcfg | show | O |
| ldapcfg | unmaprole | OM |
| passwd | OperandPresent | OM |
| passwdcfg | disableadminlockout | OM |
| passwdcfg | enableadminlockout | OM |
| passwdcfg | help | O |
| passwdcfg | set | OM |
| passwdcfg | setdefault | OM |
| passwdcfg | showall | O |
| userconfig | add | OM |
| userconfig | addad | OM |
| userconfig | addlf | OM |
| userconfig | change | OM |
| userconfig | delete | OM |
| userconfig | deletead | OM |
| userconfig | deletelf | OM |
| userconfig | showad | O |
| userconfig | showlf | O |
| userconfig | showuser | O |
| userrename | | OM |

To display the RBAC permissions for the UserManagement class:

```
switch:admin> classconfig --showroles UserManagement
Roles that have access to the RBAC class 'usermanagement' are:
```

| Role Name | Permission |
|---------------|------------|
| Admin | OM |
| Factory | OM |
| Root | OM |
| SecurityAdmin | OM |

To display the RBAC permissions for a command:

```
switch:admin> classconfig --showcli classconfig
```

| CLI | Option | Permission | RBAC Class | Context |
|-------------|-----------|------------|------------|---------|
| classconfig | help | 0 | RoleConfig | chassis |
| classconfig | show | 0 | RoleConfig | chassis |
| classconfig | showcli | 0 | RoleConfig | chassis |
| classconfig | showroles | 0 | RoleConfig | chassis |

To display the RBAC permissions for a command (in the following example, the command option entries are split and displayed in two lines because the length of the rows is too long):

```
switch:admin> classconfig --showcli ag
```

| CLI | Option | Permission | RBAC Class |
|-------|---------------------------|------------|------------|
| ag | addwwnfailov ermapping | OM | AG |
| ag | addwwnmappin g | OM | AG |
| ag | addwwnpgmapp ing | OM | AG |
| [...] | | | |

SEE ALSO [roleConfig](#)

cliHistory

Displays switch command history.

Synopsis

clihistory

clihistory --show

clihistory --showuser *username*

clihistory --showall

clihistory --clear

clihistory --help

Description

This command saves the following information whenever a command is executed on the switch:

- Timestamp
- Username
- FID (VF mode only)
- IP address of the Telnet session
- Options
- Arguments

This command displays the local CLI command history. The information is saved in the SSHOW_SYS file as part of **supportSave**. The CLI history is saved persistently to compact flash. The maximum number of saved entries for this command is 1680. CLI history records is wrapped after reaching the maximum limit specified.

The CLI history is wrapped and saved in a file that is persistent across reboots and firmware download.

Only the command name is stored in the CLI history for commands that requires password; no arguments are stored.

This command can also be executed on the standby CP.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operand:

--show

Displays the CLI history of the current user.

--showuser *username*

Displays the CLI history of the given user.

--showall

Displays the CLI history of all users.

--clear

Clears the CLI history.

--help

Displays the command usage.

Examples

To display the command history on a switch:

```

switch:admin> clihistory
Date & Time                Message
Fri Sep 19 09:43:53 2014    root, FID 10, console, tracedump
Fri Sep 19 09:43:59 2014    root, FID 10, console, coreshow
Fri Sep 19 09:44:21 2014    root, , console, firmwareshow
Fri Sep 19 09:44:25 2014    root, , console, errdump
switch:admin> firmwaredownload -s -p scp 10.70.4.109,fvt./dist.pray4green
Server IP: 10.70.4.109, Protocol IPv4
Checking system settings for firmwaredownload...
Failed to access scp://fvt:*****@10.70.4.109//dist/release.plist
switch:admin> clihistory
Date & Time                Message
Wed May 23 03:39:37 2012    admin, console, firmwaredownload

```

SEE ALSO

None

configDefault

Resets the nonpersistent configuration data to factory defaults.

Synopsis `configdefault [-fid FID | -all | -switch] [-force]`

Description Use this command to reset the nonpersistent configuration settings to their factory default values.

Configuration data is grouped into chassis information and switch information. Each configuration type is managed separately. The behavior of **configDefault** depends on the environment in which the command is executed:

- In a Virtual Fabric environment, when executed without operands, this command resets the switch configuration to default values on the current logical switch only. An Admin with chassis permissions can use additional parameters to reset configuration data for a specified logical switch (**-fid *FID***) or for all logical switches and the chassis (**-all**). Resetting default FCR configuration data requires base switch to be configured in the chassis.
- In a non-Virtual Fabric environment, when executed without operands, this command resets the switch configuration. When executed with the **-all** operand, **configDefault** resets all of the system's configuration data, including chassis and switch configurations. The **-switch** option resets the switch configuration only. The **-fid** option is not valid.

This command resets nonpersistent configuration parameters only. The following parameters are not affected by this command:

- Ethernet MAC address, IP address, subnet mask, and boot ROM parameters
- IP gateway address
- License keys
- OEM customization
- Product ID and Vendor ID
- SNMP configuration
- System name
- Chassis name
- World wide name
- Zoning configuration (includes aliases, zones, and configurations)
- Security parameters and policies
- User account passwords (includes all user configuration and all built-in accounts)
- Switch PID format
- Ethernet Link Mode

Refer to the help files for **configure** and **configureChassis** help for more information on default values for configuration parameters.

Notes This command cannot be executed on an enabled switch. You must first disable the switch using **switchDisable** or **chassisDisable**.

Some configuration parameters are cached by the system. To avoid unexpected system behavior, reboot the system after executing **configDefault**.

Note that **configDefault** does not completely remove all FCIP tunnels and GbE IP address information. This may be an issue when attempting to use the same information to create new tunnels or modify the existing ones.

The execution of this command is subject to Virtual Fabric restrictions may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

-fid *FID*

Specifies the Fabric ID of the logical switch for which to reset the configuration. This operand is valid only in Virtual Fabric mode, and the executing account must have chassis permissions.

-all

Restores all nonpersistent configuration parameters to factory default values (switch and chassis parameters). This command requires chassis permissions.

-switch

Resets the switch configuration only. This operand is not valid in VF mode.

-force

Executes the command without confirmation and reboots the switch. This operand is optional.

Examples

To restore all system configuration parameters to default values:

```
switch:admin> configdefault -all
WARNING: This is a disruptive operation that requires a switch reboot.
Would you like to continue [Y/N]: y
Executing configdefault...Please wait
Already in Native Mode continuing with configDefault
Committing configuration ... done.
The switch has to be rebooted to allow the changes to take effect.
Switch rebooting .....
Restarting system.
The system is coming up, please wait...
```

To restore all system configuration parameters with **-force** option:

```
switch:admin> configdefault -all -force
Executing configdefault...Please wait
setting the default config on switch sw_6510_79
Already in Native Mode continuing with configDefault
Committing configuration ... done.

The switch has to be rebooted to allow the changes to take effect.
Switch rebooting .....
Restarting system.
The system is coming up, please wait...
```

SEE ALSO [snmpConfig](#), [configure](#), [configureChassis](#), [switchDisable](#), [switchEnable](#)

configDownload

Downloads configuration data to the system.

Synopsis

`configdownload`

```
configdownload [- all ] [-p ftp | -ftp ] ["host", "user", "path"
["passwd"]]
```

```
configdownload [- all ] [-p scp | -scp ] [-cra] ["host", "user", "path"]
```

```
configdownload [- all ] [-p sftp | -sftp ] ["host", "user", "path"]
```

```
configdownload [- all ] [-local | -USB | -U ["file"]]
```

```
configdownload [-fid FID | -chassis | -all | -switch]
[-p ftp | -ftp ] ["host", "user", "path" [, "passwd"]]
```

```
configdownload [-fid FID | -chassis | -all | -switch ]
[-p scp | -scp ] [-cra] ["host", "user", "path"]
```

```
configdownload [-fid FID | -chassis | -all | -switch ]
[-p sftp | -sftp ] ["host", "user", "path"]
```

```
configdownload [-fid FID | -chassis | -all | -switch]
[-force] [-local | -USB | -U ["file"]]
```

```
configdownload [-vf] [-p ftp | -ftp ] ["host", "user", "path"
["passwd"]]
```

```
configdownload [ -vf] [-p scp | -scp ] [-cra] ["host", "user", "path"]
```

```
configdownload [ -vf] [-p sftp | -sftp ] ["host", "user", "path"]
```

```
configdownload [-vf] [-local | -USB | -U ["file"]]
```

```
configdownload [-fid FID | -all] -map [-p ftp | -ftp ]
["host", "user", "path" [, "passwd"]]
```

```
configdownload [-fid FID | -all] -map [-p scp | -scp ]
[-cra] ["host", "user", "path"]
```

```
configdownload [-fid FID | -all] -map [-p sftp | -sftp ]
[-cra] ["host", "user", "path"]
```

Description

This command downloads configuration parameters to the local system. Two types of configuration files can be downloaded with this command: Virtual Fabric configuration parameters and system configuration parameters. You must download both types of configuration data for the system to behave as expected. You can use FTP or SCP to download configuration files from a remote host, or you can retrieve the configuration files from a predetermined directory on the local system, or from an attached USB device.

Use the **-vf** option to download the Virtual Fabric configuration data. The Virtual Fabric configuration file includes logical switch definitions and Virtual Fabric status (enabled or disabled). The file should be named `vf-conf_xx.txt` to distinguish it from the regular system configuration (`config.txt`). The `xx` indicates the platform ID. Virtual Fabric configuration data can only be shared between switches that belong to the same platform type and share the same platform ID. If the platform ID contained in the header of the configuration file does not match the platform ID of the system to which it is downloaded, **configDownload** fails. When you download a `vf-conf_xx.txt` file, all attributes defined in this file are downloaded to the system and take effect with the exception of LISL ports. The LISL ports on the system are not affected by this download.

Use the **-all** option to download the system configuration data (which must be downloaded separately from the Virtual Fabric configuration data). It is grouped into chassis information and switch information. Each configuration type is managed separately and the behavior of **configDownload** depends on the environment in which the command is executed and which part of the system configuration you wish to download.

- In a Virtual Fabric environment, when the **configDownload** command is executed without chassis permissions, this command downloads the switch configuration to the current logical switch only. An Admin user with chassis permissions can use additional parameters to perform the following selective configuration downloads:
 - Download the switch configuration to a specified logical switch (**-fid FID**).
 - Download the chassis configuration only (**-chassis**).
 - Download the entire configuration including the data for all logical switches and for the chassis (**-all**).

The interactive version of the command (no operands) prompts for input on only the parameters the user is allowed to execute.
- In a non-Virtual Fabric environment, this command by default downloads the configuration for the default logical switch only. To download the chassis-level configuration only, use the **-chassis** option. To download both the chassis and switch configuration, use the **-all** option. Chassis permissions are required. The **-fid** option is not valid. The **-switch** option is equivalent to issuing the command default (without options).

Perform the following steps to backup and then restore a configuration in a switch using Virtual Fabrics:

1. Run the **configUpload -vf** command followed by the **configUpload -all** command from the old setup.
2. Run the **configDownload -vf** command followed by the **configDownload -all** command in the new setup.

The switch must be disabled for configuration download of all parameters (with the exception of SNMP and MAPS).

The following rules apply to configuration download in Virtual Fabric mode:

- When downloading the chassis configuration, the number of logical switches defined in the configuration download must match the number of logical switches currently defined on the switch.
- When downloading the switch configuration, the target FID must be defined in both the configuration download and the current system.
- When downloading the switch configuration from a specified source FID to a target FID, the target FID must be defined on the switch and the source FID and associated configuration must be defined in the configuration download. In addition, downloading an SFID configuration resets the target FID ports without warning. Caution is advised when using this option.
- When downloading all configuration parameters, the number of switches defined in the downloaded configuration file must match the number of switches currently defined on the switch. In addition, the following restrictions apply:
 - The switches must be disabled unless you only wish to download SNMP or MAPS parameters.
 - Downloading a configuration file from a system that is not Virtual Fabric-capable to a system in Virtual Fabric mode is not recommended. The configuration is applied to the default switch only, and only to the ports that are part of the default switch.

If an FCS policy is enabled, the following rules and restrictions apply:

- Both [Defined Security Policies] and [Active Security Policies] sections must exist and contain the FCS_POLICY.
- In the [Defined Security Policies] section, at least one member of the FCS_POLICY must be the same as a member in the previous FCS_POLICY.
- In the [Active Security Policies] section, the FCS_POLICY must be exactly the same as the previous FCS_POLICY. Order of members must be maintained.
- If either security policies section has an RSNMP_POLICY, then that section must have a WSNMP_POLICY.
- After the switch is enabled, if the switch is the primary FCS, then its security and zoning information is propagated to all other switches in the fabric.
- After the switch is enabled, if the switch is a non-FCS or a backup FCS, then its security and zoning information will be overwritten by the primary FCS.

Security parameters and the switch identity cannot be changed by **configDownload**. Parameters such as the switch name and IP address are ignored; they are lines in the configuration file that begin with "boot". Security parameters and version stamp are ignored; they are the lines in the configuration file that begin with "sec".

[License] is only accepted if the boot.mac parameter matches the license ID (WWN) of the switch performing the download; otherwise, it is ignored.

The configuration parameters R_A_TOV, E_D_TOV, WAN_TOV, and MAX_HOPS are interrelated. Assigning a specific value to one or more of these parameters might change the range of allowed values that can be assigned to the other parameters. As a result, you may not be able to set all the values within the range displayed for each parameter. This command validates the modified values of these four parameters and terminates the download operation, if the validation check fails.

This is particularly important when downloading a zoning configuration. Because the new zoning information is added to the current configuration, there might not be any conflicts. If the current zoning configuration is to be replaced, the keyword "clear:" should be inserted into the configuration file immediately before the zoning lines (starting at the line "[Zoning]").

configDownload

If the configuration file contains the keyword "enable:" followed by a *zone_configuration*, that zoning configuration is enabled in the fabric. If there is no "enable:" keyword in the configuration file or no zoning configuration by that name exists, or if enable fails for any reason (such as dangling aliases), then the following conditions apply:

- The effective configuration remains as it was prior to the configuration download. The "enable:" action is ignored.
- The Defined Configuration changes to reflect the new zoning configuration.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

If the switch is in Access Gateway mode, some of the configuration parameters will take effect only after the switch has been re-enabled with the **switchEnable** command.

F_Port trunking configurations are not saved in the configuration file and will not be restored after a configuration download.

Do not manually edit a configuration file after uploading the file and before downloading the file to a switch. Manual editing bypasses sanity checks for some configuration parameters and results in unpredictable system behavior.

The **configDownload** command prompts for confirmation and issues a notice that "A switch reboot is required for the changes to take effect." A configuration download affects large amount of configuration data, and there is no reliable mechanism to determine, which of these parameters may or may not require a reboot. To ensure that all configuration changes are applied correctly, Brocade strongly recommends that you always reboot the switch after a configuration download.

Operands This command has the following operands:

-p ftp | -ftp or -p scp | -scp or -p sftp | -sftp

Specifies the data transmission protocol as either file transfer protocol (FTP), secure copy protocol (SCP), or secure FTP (SFTP). If no protocol is specified, the protocol defaults to FTP.

-cra

Enables Challenge Response Authentication (CRA). CRA is supported only with the SCP protocol.

-vf

Downloads the Virtual Fabric configuration (vf-conf_xx.txt) instead of the regular system configuration. The vf-con_xx.txt file contains a listing of logical switches configured on the platform specified by the platform ID (xx) and other Virtual Fabric parameters. You cannot use the **-vf** option with any of the system configuration upload options (**-fid**, **-chassis**, **-all**).

-all

Downloads all configuration data, including chassis and switch configuration data.

-fid FID

Downloads the switch configuration to a logical switch specified by its fabric ID. This operand is valid only in a Virtual Fabric environment and requires chassis permissions.

-chassis

Downloads the chassis configuration only.

| | |
|------------------|--|
| -switch | Downloads the switch configuration only. This operand is valid only in non-VF mode. |
| -map | Downloads the port-to-area addressing mode configuration files. |
| "host" | Specifies the name or the IP address of the external host, from which to download the configuration. IPv4 and IPv6 addresses are supported. To be able to mention the FTP server by name, you need to set up two DNS servers with dnsConfig . Quotation marks are optional. |
| "user" | Specifies the login name for the external host. Quotation marks are optional. |
| "path" | Specifies the file name and path of the configuration file. When used with the -map option, this parameter specifies a folder that contains all port-to-area mapping files. Absolute path names may be specified using a forward slash (/). Relative path names search for the file in the login account's home directory on UNIX hosts and in the directory on which the FTP server is running on Windows hosts. This operand is valid only when the file is downloaded from an external host. Quotation marks are optional. |
| "passwd" | Specifies the account password. Quotation marks are optional. |
| -local | Downloads a specified configuration file from a predetermined directory on the local chassis. |
| -USB -U | Downloads a specified configuration file from a predetermined directory on an attached USB device. |
| "file" | Specifies the name of the configuration file to be downloaded. This parameter can be used only with the -local or -USB option, each of which retrieves files from a predetermined directory on the local chassis or on an attached USB device. Therefore, subdirectories and absolute path names are not permitted. Quotation marks are optional. |
| -force | Overwrites an existing file without confirmation. This parameter is valid only with the -local or -USB options. |

Diagnostics

The configuration download may fail for one or more of the following reasons:

- The switch has not been disabled. Disabling the switch is not necessary for configuration files containing only certain SNMP or MAPS parameters. You can run **configDownload** first without disabling the switch, and if there is at least one changed parameter outside of SNMP or MAPS, you are prompted to disable the switch before proceeding.
- The host name is not known to the switch.
- The host IP address cannot be contacted.
- You do not have permission on the host.
- You are running a script that prints something at login.
- The file does not exist on the host.
- The file is not a switch configuration file.
- The FTP server is not running on the host.
- The configuration file contains errors.
- The configuration file's logical switch definitions do not match the definitions on the target switch.

Examples

To download the switch configuration file interactively to the current logical switch from a local directory (no chassis permissions):

```
switch:admin> configdownload
Protocol (scp, ftp, sftp, local) [ftp]:
Server Name or IP Address [host]: 192.168.163.233
User Name [user]: admin
Path/Filename [<home dir>/config.txt]:
Section (all|chassis|FID# [all]):
      *** CAUTION ***
```

This command is used to download a backed-up configuration for a specific switch. If using a file from a different switch, this file's configuration settings will override any current switch settings. Downloading a configuration file, which was uploaded from a different type of switch, may cause the switch to fail.

A switch reboot is required for the changes to take effect.

Please make sure all the switches are disabled by using `chassisdisable` command. Downloading configuration to an online switch may result in some configuration not being downloaded to that switch.

```
configDownload operation may take several minutes
to complete for large files.
Do you want to continue [y/n]: y
Password: *****
Activating configDownload: Switch is disabled
configDownload complete: All config parameters are downloaded
```

To download the switch configuration data to the current logical switch from an external FTP server (no chassis permissions):

```
switch:admin> configdownload -ftp 192.168.38.245, \
      jdoe.config.txt,password
```

To download all system configuration data for the chassis and all logical switches (requires chassis permissions):

```
switch:admin> configdownload -all -ftp 192.168.38.245, \  
jdoe,config.txt,password
```

To download the switch configurations to a logical switch with FID 8 from an attached USB device (requires chassis permissions):

```
switch:admin> configdownload -fid 8 -USB config.txt
```

To download the switch configurations belonging to a logical switch with FID 4 to a logical switch with FID 8 from an attached USB device (requires chassis permissions):

```
switch:admin> configdownload -fid 8 -sfid 4 \  
-USB config_fid8.txt
```

To download the Virtual Fabric configuration file using secure FTP:

```
switch:admin> configdownload -vf -p sftp \  
10.32.248.119,jdoe,/temp/vf-conf_66.txt,password
```

SEE ALSO [configDefault](#), [configList](#), [configShow](#), [configUpload](#), [configure](#), [configRemove](#)

configList

Lists uploaded configuration files.

Synopsis `configlist -local | -USB | -U`

Description This command displays a list of names, sizes, and creation dates of configuration files saved on the local chassis or on an attached USB device. These files are created when the **configUpload** command is executed with the **-local** or the **-USB** option.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-local

Displays the list of configuration files on the local chassis.

-USB | -U

Displays the list of configuration files on the attached USB device.

Examples To display a list of configuration files stored on the local chassis:

```
switch:admin> configlist -local
config.txt      25679      2007 Jan 02 15:16
config2.txt     25679      2007 Jan 06 15:16
next_cfg.txt    20977      2007 Jan 18 15:16
```

SEE ALSO [configDownload](#), [configUpload](#), [configShow](#), [configRemove](#)

configRemove

Deletes a saved configuration file.

Synopsis `configremove -local | -USB | -U [file]`

Description This command deletes a configuration file that was previously saved to the local chassis or to an attached USB device by using the **configUpload** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-local Removes a configuration file that was previously created by **configUpload -local** from the local chassis.

-USB | -U Removes a configuration file that was previously created by **configUpload -USB** from an attached USB device.

file Specifies the configuration file to be removed. If the **file** option is omitted, the command prompts for a file name.

Examples To remove a configuration file from the local chassis:

```
switch:admin> configremove -local first_config.txt
```

To remove a configuration file from an attached USB device without specifying a filename:

```
switch:admin> configremove -USB
File Name [config.txt]: second_config.txt
```

SEE ALSO [configDownload](#), [configUpload](#), [configList](#), [configShow](#)

configShow

Displays system configuration settings.

Synopsis

configshow

```
configshow [-all | -fid FID | -chassis | -switch] |
  [-local | -USB | -U] [file] [-pattern "pattern"]
```

Description

Use this command to display system configuration settings. Some but not all of these parameters are set by the **configure** and **configureChassis** commands.

Configuration data is grouped into chassis information and switch information. Each configuration type is managed separately. The behavior of **configShow** depends on the environment in which the command is executed:

- In a Virtual Fabric environment, when executed without operands, this command displays the switch configuration for the current logical switch. An Admin with chassis permissions can use additional parameters to display configuration data for a specified logical switch (**-fid *FID***), for the chassis (**-chassis**), or for all logical switches and the chassis (**-all**).
- In a non-Virtual Fabric environment, when executed without operands, this command displays the switch configuration. When executed with the **-all** operand, **configShow** displays all of the system's configuration data, including chassis and switch configuration data. The **-chassis** option displays the chassis configuration only. The **-switch** option displays the switch configuration only. The **-fid** option is not valid.

Notes

Not all values displayed are applicable to all system models and configurations.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

The following operands are optional:

-pattern "*pattern*"

Specifies a text string, enclosed in double quotation marks, that limits the output of the command to only those entries that contain the pattern. Use of wildcards and other common regular expression operators is not supported. Some configuration settings do not display when filtered. When in doubt, use the command without the **-pattern** operand. Executing **configShow -pattern** without further operands is not recommended and can cause unspecified behavior.

-all

Displays all configuration data including chassis and switch configuration.

-fid *FID*

Displays configuration data for a logical switch specified by its fabric ID. This option is valid only in a Virtual Fabric environment and requires chassis permissions

-chassis

Displays configuration data for the chassis only. This option is valid only in a Virtual Fabric environment and requires chassis permissions

-switch

Displays the switch configuration only. This option is valid only in a Non-Virtual Fabric environment.

-local [file]

Displays the content of a configuration file that was previously created by **configUpload** and stored on the chassis. The output can be optionally filtered by **-pattern "pattern"**. If **file** is omitted, the command prompts for a file name. The output format when **-local** is specified matches that of **configUpload** and contains a superset of the information provided when **-local** is not specified.

-USB | -U [file]

Displays the content of a configuration file that was previously created by **configUpload** and stored on an attached USB device. The output can be optionally filtered by **-pattern "pattern"**. If **file** is omitted, the command prompts for a file name. The output format when **-USB** is specified matches that of **configUpload** and contains a superset of the information provided when **-USB** is not specified.

Examples

To display all configuration data on a Virtual Fabric-enabled system:

```
switch :admin> configshow -all
[Configuration upload Information]
Configuration Format = 3.0
Minimum Compatible Format = 3.0
Excluding Format = 0.0
date = Sun May 26 06:18:40 2013
FOS version = v7.2.0
Number of LS = 2
[Chassis Configuration Begin]

[fcRouting]
fcRoute.backboneFabricId:100
fcRoute.fcrState:2
fcRouteParam.maxLsanCount:3000
fcRoute.port.8.xportAdmin:DISABLED
fcRoute.port.8.fabricId:4
fcRoute.port.8.ratov:10000
fcRoute.port.8.edtov:2000
fcRoute.port.8.frontConfigDid:160
fcRoute.port.8.portType:400
fcRoute.port.8.portMode:0
fcRoute.port.8.autoElp:7
fcRoute.port.9.xportAdmin:DISABLED
fcRoute.port.9.fabricId:5
fcRoute.port.9.ratov:10000
fcRoute.port.9.edtov:2000
fcRoute.port.9.frontConfigDid:160
fcRoute.port.9.portType:400
fcRoute.port.9.portMode:0
fcRoute.port.9.autoElp:7
fcRouteParam.port.8.rportCost:0
fcRouteParam.port.9.rportCost:0
fcRoute.xlate.persistxdState:1
fcRouteParam.lsan.tagCnt:0

[Chassis Configuration]
passwdcfg.minlength:8
passwdcfg.lowercase:0
passwdcfg.uppercase:0
passwdcfg.digits:0
passwdcfg.punctuation:0
```

configShow

```
passwdcfg.history:1
passwdcfg.minpasswordage:0
passwdcfg.maxpasswordage:0
passwdcfg.warning:0
passwdcfg.lockoutthreshold:0
passwdcfg.lockoutduration:30
passwdcfg.adminlockout:0
passwdcfg.repeat:1
passwdcfg.sequence:1
passwdcfg.status:0
fips.mode:0
fips.selftests:0
ipfilter.0.name:default_ipv4
ipfilter.0.numofrules:12
ipfilter.0.rule.1:0,0x23,0,0,6,22
ipfilter.0.rule.10:0,0x23,0,0,17,123
ipfilter.0.rule.11:0,0x63,0,0,6,600,1023
ipfilter.0.rule.12:0,0x63,0,0,17,600,1023
ipfilter.0.rule.2:0,0x23,0,0,6,23
ipfilter.0.rule.3:0,0x23,0,0,6,897
ipfilter.0.rule.4:0,0x23,0,0,6,898
ipfilter.0.rule.5:0,0x23,0,0,6,111
ipfilter.0.rule.6:0,0x23,0,0,6,80
ipfilter.0.rule.7:0,0x23,0,0,6,443
ipfilter.0.rule.8:0,0x23,0,0,17,161
ipfilter.0.rule.9:0,0x23,0,0,17,111
ipfilter.0.state:3
ipfilter.0.type:0
ipfilter.1.name:default_ipv6
ipfilter.1.numofrules:12
ipfilter.1.rule.1:0,0x23,0,0,6,22
ipfilter.1.rule.10:0,0x23,0,0,17,123
ipfilter.1.rule.11:0,0x63,0,0,6,600,1023
ipfilter.1.rule.12:0,0x63,0,0,17,600,1023
ipfilter.1.rule.2:0,0x23,0,0,6,23
ipfilter.1.rule.3:0,0x23,0,0,6,897
ipfilter.1.rule.4:0,0x23,0,0,6,898
ipfilter.1.rule.5:0,0x23,0,0,6,111
ipfilter.1.rule.6:0,0x23,0,0,6,80
ipfilter.1.rule.7:0,0x23,0,0,6,443
ipfilter.1.rule.8:0,0x23,0,0,17,161
ipfilter.1.rule.9:0,0x23,0,0,17,111
ipfilter.1.state:3
ipfilter.1.type:1
(output truncated)
```

To filter the content to display only the password configuration:

```
switch :admin> configshow -all -pattern "passwdcfg"
passwdcfg.minlength:8
passwdcfg.lowercase:0
passwdcfg.uppercase:0
passwdcfg.digits:0
passwdcfg.punctuation:0
passwdcfg.history:1
passwdcfg.minpasswordage:0
passwdcfg.maxpasswordage:0
passwdcfg.warning:0
passwdcfg.lockoutthreshold:0
```

```

passwdcfg.lockoutduration:30
passwdcfg.adminlockout:0
passwdcfg.repeat:1
passwdcfg.sequence:1
passwdcfg.status:0

```

To display switch configuration data for FID 20:

```

switch :admin> configshow -fid 20
[Configuration upload Information]
Configuration Format = 2.0
date = Tue Oct 7 14:53:12 2008
FOS version = v6.2.0.0
Number of LS = 3
[Switch Configuration Begin : 0]
SwitchName = Spirit_66
Fabric ID = 20

```

```

[Boot Parameters]
boot.name: Spirit_66
boot.ipa: 10.32.228.66
boot.lidid: 10:00:00:05:1e:41:5c:c1
boot.mac: 10:00:00:05:1e:41:5c:c1
boot.device: eth0
boot.gateway.ipa: 10.32.224.1

```

```

[Configuration]
acl.clear:0
ag.port.nfportfailback:0x0
ag.port.nfportfailover:0x0
ag.port.nfporttopo.0:0x0
ag.port.nfporttopo.1:0x0
ag.port.nfporttopo.10:0x0
ag.port.nfporttopo.11:0x0
ag.port.nfporttopo.12:0x0
ag.port.nfporttopo.13:0x0
ag.port.nfporttopo.14:0x0
ag.port.nfporttopo.15:0x0
ag.port.nfporttopo.16:0x0
ag.port.nfporttopo.17:0x0
ag.port.nfporttopo.18:0x0
ag.port.nfporttopo.19:0x0
ag.port.nfporttopo.2:0x0
ag.port.nfporttopo.20:0x0
(output truncated)

```

SEE ALSO [configure](#), [configureChassis](#), [configDownload](#), [configUpload](#), [configList](#), [configRemove](#), [diagDisablePost](#), [diagEnablePost](#), [ipAddrShow](#), [licenseShow](#)

configUpload

Uploads system configuration data to a file.

Synopsis

configupload

```
configupload [-all] [-p ftp | -ftp] ["host", "user", "path"
  [, "passwd"]]
```

```
configupload [-all] [-p scp | -scp] [-cra] ["host", "user", "path"]
```

```
configupload [-all] [-p sftp | -sftp] ["host", "user", "path"]
```

```
configupload [-all] [-force] [-local | USB |-U] ["file"]
```

```
configupload [-fid FID | -chassis | -all | -switch]
  [-p ftp | -ftp] ["host", "user", "path" [, "passwd"]]
```

```
configupload [-fid FID | -chassis | -all | -switch]
  [-p scp | -scp] [-cra] ["host", "user", "path"]
```

```
configupload [-fid FID | -chassis | -all | -switch]
  [-p sftp | -sftp] ["host", "user", "path"]
```

```
configupload [-fid FID | -chassis | -all] | -switch]
  [-force] [-local | USB |-U] ["file"]
```

```
configupload [-vf] [-p ftp | -ftp] ["host", "user", "path"
  [, "passwd"]]
```

```
configupload [-vf] [-p scp | -scp] [-cra] ["host", "user", "path"]
```

```
configupload [-vf] [-p sftp | -sftp] ["host", "user", "path"]
```

```
configupload [-vf] [-force] [-local | USB |-U] ["file"]
```

```
configupload [-fid FID | -all] -map [-p ftp | -ftp ]
  ["host", "user", "path" [, "passwd"]]
```

```
configupload [-fid FID | -all] -map [-p scp | -scp ]
  [-cra] ["host", "user", "path"]
```

```
configupload [-fid FID | -all] -map [-p sftp | -sftp ]
  [-cra] ["host", "user", "path"]
```

Description

This command uploads configuration data to a file. Two types of configuration files can be uploaded with this command: Virtual Fabric configuration parameters and system configuration parameters.

Use the **-vf** option to upload Virtual Fabric configuration parameters. The Virtual Fabric configuration includes logical switch definitions and Virtual Fabric status (enabled or disabled). The file should be named `vf-conf_xx.txt` to distinguish it from the regular system configuration (`config.txt`). The `xx` indicates the platform ID specified in the header of the configuration file. The platform ID is the same as the first two digits of the "switchType" parameter displayed by **switchShow**. Virtual Fabric configuration data can only be shared between switches that belong to the same platform type and share the same platform ID. Refer to **configDownload** help for more information on the Virtual Fabric configuration.

The system configuration data is uploaded separately. It is grouped into chassis information and switch information. Each configuration type is managed separately and the behavior of **configUpload** depends on the environment in which the command is executed and which part of the system configuration you wish to upload.

- In a Virtual Fabric environment, when executed without chassis permissions, this command uploads the current logical switch configuration only. An Admin user with chassis permissions can use additional parameters to perform the following selective configuration uploads:
 - Upload the switch configuration of a specified logical switch (**-fid FID**).
 - Upload the chassis configuration only (**-chassis**).
 - Upload the entire system configuration including the data for all logical switches and for the chassis (**-all**).
 - Upload the switch configuration only in Non-VF mode (**-switch**).

The interactive version of the command (no operands) prompts for input on only the parameters the user is allowed to execute.

- In a non-Virtual Fabric environment, this command by default uploads the configuration for the default logical switch only. To upload the chassis-level configuration only, use the **-chassis** option. To upload both the chassis and switch configuration, use the **-all** option. Chassis permissions are required. The **-fid** option is not valid. The **-switch** option is equivalent to issuing the command default (without options).

Perform the following steps to backup and then restore a configuration in a switch using Virtual Fabrics:

1. Run the **configUpload -vf** command followed by the **configUpload -all** command from the old setup.
2. Run the **configDownload -vf** command followed by the **configDownload -all** command in the new setup.

You can use the file transfer protocol (FTP), the secure copy protocol (SCP), or secure FTP (SFTP) to upload configuration files to an external host, or you can save the configuration in a predetermined directory on the local chassis or on an attached USB device. If the specified file already exists, this command prompts you to overwrite the file. Specify **-force** to overwrite the file without confirmation. When the local chassis is chosen as the destination, the resulting file is written to both primary and secondary partitions, and on enterprise-class platforms, to both Active and Standby Control Processors (CPs).

Refer to the *Fabric OS Administrator's Guide* for information on backward compatibility and on the content of the configuration file.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

F_Port trunking configurations are not saved in the configuration file and will not be restored after a configuration download.

Do not manually edit a configuration or a vf-conf.xx file after uploading the file and before downloading the file to a switch. Manual editing bypasses sanity checks for some configuration parameters and results in unpredictable system behavior.

Operands

This command has the following operands:

- p ftp | -ftp or -p scp | -scp or -p sftp | -sftp**
Specifies the data transmission protocol as either file transfer protocol (FTP), secure copy protocol (SCP), or secure FTP (SFTP). If no protocol is specified, the protocol defaults to FTP.
- cra**
Enables Challenge Response Authentication (CRA). CRA is supported only with the SCP protocol.
- vf**
Uploads the Virtual fabric configuration to a file. You must specify a filename when uploading this file. It is recommended to name this file vf-conf_xx.txt (where xx indicates the platform ID) to distinguish this file from the system configuration (config.txt). Use **switchShow** to determine the platform ID of the system. The platform ID in the header of the configuration file is the same as the first two digits of the switchType parameter in the **switchShow** output. You cannot use the **-vf** option with any of the regular configuration upload options (**-fid**, **-chassis**, **-all**).
- fid FID**
Uploads switch configuration data from a logical switch specified by its fabric ID. This parameter is valid only in a Virtual Fabric environment and requires chassis permissions.
- chassis**
Uploads chassis configuration only.
- all**
Uploads all system configuration data including chassis and switch configuration for all logical switches.
- switch**
Uploads the switch configuration only. This operand is not valid in VF mode.
- map**
Uploads the port-to-area addressing mode configuration files.
- "host"**
Specifies the name or the IP address of the external host to which to upload the configuration. To be able to mention the FTP server by name, you need to set up one or more DNS servers with **dnsConfig**. Quotation marks are optional.
- "user"**
Specifies the login name for the external host. Quotation marks are optional.

| | |
|------------------|--|
| "path" | Specifies the file name and path of the configuration file. When used with the -map option, this parameter specifies a folder that contains all port-to-area mapping files. Absolute path names may be specified using a forward slash (/). Relative path names upload the file to the login account's home directory on UNIX hosts and into the directory on which the FTP server is running on Windows hosts. This operand is valid only when the file is uploaded to an external host. Quotation marks are optional. |
| "passwd" | Specifies the account password. Quotation marks are optional. |
| -local | Uploads a specified configuration file to a predetermined directory on the local chassis. This option requires a file name. |
| -USB -U | Uploads a specified configuration file to a predetermined directory on an attached USB device. This option requires a file name. |
| "file" | Specifies the file name. Quotation marks are optional. This parameter is valid only with the -local or -USB options, each of which stores files in a predetermined directory on the local chassis or on an attached USB device. Therefore, subdirectories and absolute path names are not permitted. |
| -force | Overwrites an existing file without confirmation. This parameter is valid only with the -local or -USB options. |

When invoked without operands or without **"host"** or **"file"** parameters, **configUpload** runs in interactive mode.

Diagnostics The configuration upload may fail for one or more of the following reasons:

- The host name is not known to the switch.
- The host IP address cannot be contacted.
- The user does not have permission on the host.
- The FTP server is not running on the host.

Examples To upload the switch configuration interactively from a switch that is not enabled for Virtual Fabrics:

```
switch:admin> configupload
Protocol (scp, ftp, sftp, local) [ftp]: sftp
Server Name or IP Address [host]: 192.168.38.245
User Name [user]: jdoe
File Name [<home dir>/config.txt]:
Section (all|chassis|switch [all]):
Password: *****

configUpload complete: All config parameters are uploaded
```

configUpload

To upload the switch configuration that belongs to a logical switch with FID 100:

```
switch:admin> configupload
Protocol (scp, ftp, sftp, local) [ftp]: sftp
Server Name or IP Address [host]: 10.32.220.100
User Name [user]: jdoe
File Name [<home dir>/config.txt]: config.fid100.txt
Section (all|chassis|FID# [all]): 100
Password: *****
```

```
configUpload complete: All config parameters are uploaded
```

To upload the configuration for the entire chassis to a local file from the command line forcing an overwrite:

```
switch:admin> configupload -chassis -local \
-force config.txt
```

```
configUpload complete: All config parameters are uploaded
```

To upload the configuration for the current logical switch to an external FTP server:

```
switch:admin> configupload -ftp 192.168.38.245, \
jdoe,config.txt,password
```

To upload all system configuration data to an external FTP server:

```
switch:admin> configupload -all -ftp \
192.168.38.245,jdoe,config.txt,password
```

To upload the system configuration file for a logical switch with FID 8 to an attached USB device:

```
switch:admin> configupload -fid 8 \
-USB config.txt
```

To upload the Virtual Fabric configuration of the current platform to an external FTP server:

```
switch:admin> configupload -vf -p ftp \
10.32.248.119,jdoe,/temp/vf-conf.66.txt,password
```

SEE ALSO [configDefault](#), [configDownload](#), [configShow](#), [configList](#), [configRemove](#), [configure](#), [configureChassis](#)

configure

Changes switch configuration settings.

Synopsis **configure**

Description Use this command to change switch configuration settings.

Configuration data is grouped into chassis information and switch information. Each configuration type is managed separately. For information on file format and specific parameters contained in each section, refer to the **configUpload** help page.

The behavior of the **configure** command depends on the environment in which the command is executed:

- In a Virtual Fabric environment, the **configure** command sets switch configuration parameters for the current logical switch only. If a switch or chassis is configured with multiple logical switches, you must configure each logical switch separately. Use the **setContext** command to change the current logical switch context.
- In a non-Virtual Fabric environment, the **configure** command sets switch configuration parameters.

To configure chassis-wide parameters, use the **configureChassis** command.

The following switch configuration parameters can be set with the **configure** command:

- Switch fabric parameters
- Virtual channel parameters
- F_Port login parameters
- Zoning operation parameters
- Remote State Change Notifications (RSCN) transmission mode
- Arbitrated Loop parameters
- System Services settings
- Portlog Events enable or disable settings
- Secure Socket Layer (SSL) attributes
- Remote Procedure Call Daemon (RPCD) attributes
- Web Tools attributes

To access all parameters controlled by this command, you must disable the switch using the **switchDisable** command. If executed on an enabled switch, only a subset of attributes are configurable. Menu displays may vary depending on the hardware platform.

The **configure** command runs in interactive mode and presents you with a series of hierarchical menus. Each top-level menu and its associated submenus consist of a text prompt, a selection of valid values, and a default value (in brackets).

The following keys control the execution of the command:

Return

When entered at a prompt with no preceding input, the command accepts the default value (if applicable) and moves to the next prompt.

Interrupt (Ctrl-C)

Aborts the command immediately and ignores all changes made.

End-of-file (Ctrl-D)

When entered at a prompt with no preceding input, terminates the command and saves changes made.

The following parameters can be modified with the **configure** command:

Fabric Parameters

Fabric settings control the overall behavior and operation of the fabric. Some of these settings, such as the domain, are assigned automatically and may differ from one switch to another in a given fabric. Other parameters, such as buffer-to-buffer credit or timeout values, can be modified to suit specific applications or operating environments but must be in agreement among all switches to allow formation of the fabric.

The following fabric settings can be modified (* = multiplication symbol)

| Field | Type | Default | Range |
|----------------------------|---------|---------|--------------------------|
| Domain | Number | 1 | 1-239 |
| Enable 8-bit | | | |
| Dynamic Area Mode | Boolean | 0 | 0 or 1 |
| WWN Based persistent PID | Boolean | no | yes/no |
| D-Port Parameters | Boolean | on | on/off |
| Dynamic D-Port | Boolean | on | on/off |
| On-Demand D-Port | Boolean | on | on/off |
| RDP Polling Cycle | Number | 1 | 0 to 24 |
| Allow XISL Use | Boolean | yes | yes/no |
| Remote Fosexec | Boolean | on | on/off |
| High Integrity Fabric Mode | Boolean | on | on/off |
| Display FDMI Host Name | Boolean | on | on/off |
| Location ID | Number | 0 | 0 to 4 |
| Disable FID Check | Boolean | no | yes/no |
| Enable a 256 Area Limit | Number | 0 | 0 to 2 |
| R_A_TOV | Number | 10000 | E_D_TOV * 2 to 120000 |
| E_D_TOV | Number | 2000 | 1000 to R_A_TOV/2 |
| WAN_TOV | Number | 0 | 0 to R_A_TOV/4 |
| MAX_HOPS | Number | 7 | 7 to 19 |
| Data Field Size | Number | 2112 | 256 to 2112 |
| Sequence Level Switching | Boolean | 0 | 0 or 1 |
| Disable Device Probing | Boolean | 0 | 0 or 1 |
| Suppress Class F Traffic | Boolean | 0 | 0 or 1 |
| Per-frame Route Priority | Boolean | 0 | 0 or 1 |
| Long Distance Fabric | Boolean | 0 | 0 or 1 |
| BB Credit | Number | 16 | 1 to 27 |
| Disable FID Check | Boolean | yes | yes/no |
| Insistent Domain ID Mode | Boolean | no | yes/no |
| Disable Default PortName | Boolean | no | yes/no |
| Dynamic Portname | Boolean | no | yes/no |
| Edge hold time | Number | 220 | 80 to 500 |
| F-Port Device Update Mode | Boolean | no | yes/no |

Fabric parameters are defined as follows:

Domain

The domain number uniquely identifies a switch in a fabric. This value is automatically assigned by the fabric. The range is 1-239.

Enable 8-bit Dynamic Area Mode

When enabled, this feature supports Dynamic Area Mode in default partitions on the Brocade DCX 8510-8 and DCX 8510-4. Dynamic Area Mode is disabled by default. When enabled, Dynamic Area Mode supports both static and dynamic area assignment. Use the **portAddress** command to perform a static assignment of an area to a given port. In Dynamic Area Mode, areas are dynamically assigned to the ports (up to a 255 limit). Port area assignments are persistent; however, disabling Dynamic Area Mode with **configure** resets the area assignments. This feature is configurable only on the default switch.

Enabling Dynamic Area Mode fails under one or more of the following conditions:

- The number of ports in the default partition exceeds 255.
- An AP blade with FL ports is present in the chassis (Brocade Encryption blade, or FCoE 10-24).

WWN Based persistent PID

When enabled, this feature supports both dynamic and static WWN-based PID assignment. In dynamic PID binding, the first area assigned to a device when it logs in is bound to the device WWN and remains persistent through subsequent logins. Every time the device logs into the switch, it is guaranteed to get the same PID. Alternately, you can use the **wwnAddress** command to create a static WWN-based PID assignment. In either case, the WWN-based persistent PID feature must be enabled through **configure**. The feature is disabled by default; it is dependent on Dynamic Area Mode being enabled.

D-Port Parameters

Disables or enables the D_Port parameters on the switch.

Dynamic D_Port

Disables or enables Dynamic D_Port mode configuration on the switch. By default, Dynamic D_Port mode is ON. When Dynamic D_Port mode is ON, the port may dynamically go into D_Port mode based on external request from remote HBA or device port. After the D_Port tests are complete, the port may automatically switch to operate as an F_Port. This option is not supported in AG mode.

On Demand D_Port

Disables or enables On-Demand D_Port mode configuration on the switch. By default, On-Demand D_Port mode is OFF.

RDP Polling Cycle

Displays the Switch Driver timer routine for polling. The switch need not be disabled to configure this variable.

Allow XISL Use

An extended interswitch link (XISL) is an interswitch link (ISL) that connects the logical switch to the base switch and carries traffic for multiple logical fabrics. This feature is supported only on Virtual Fabric-aware platforms under the following conditions: Virtual Fabrics must be enabled on the switch, and the switch cannot be a base switch. In addition, on the

Brocade DCX 8510-8 and DCX 8510-4 the switch cannot be a default switch or include interchassis link (ICL) ports. This feature is enabled by default (yes=enabled). On the Brocade 6510 default switch, the feature is disabled by default (no=disabled). The XISL feature is supported on the Brocade extension blade (FX8-24).

You do not need to disable the logical switch before changing the value of this parameter. Turning off XISL use requires confirmation because all LISLs will be removed upon execution. If the logical switch is enabled and is part of an edge fabric connected to an FCR, this parameter cannot be turned on. If the logical switch is disabled or it is not yet part of an edge fabric, this parameter can be turned on. However, execution may cause edge fabric segmentation if the EX_Port connected to the edge fabric is disabled while the logical switch is enabled or connected to the edge fabric.

Location ID

The Location ID parameter affects routing calculations, and should remain set to the default value of 0 for normal use. Do not change the value unless explicitly instructed to do so by a Brocade Support engineer.

Disable FID Check

If fabric ID (FID) check is disabled, the fabric ignores the Fabric Identifier conflict with the neighboring switch during fabric formation. By default, FID check is enabled. If the fabric detects a FID conflict, it disables the E_Port with a "Fabric ID conflict" message. This parameter is configurable only if the switch is Virtual Fabric-aware and Virtual Fabric is enabled on the switch.

Enable 256 Area limit

The 256 area limit allows the partition to be configured for 8-bit addressing rather than the default 10-bit addressing. Each port in this partition is given a unique area represented by the middle 8 bits of the PID. Valid values include the following:

0

No limit is imposed on the area. This is the default value. The partition is configured for 10-bit addressing and supports up to 1800 ports.

1

The unique area assignments begin at zero regardless of where the port is physically located. This allows FICON users to make use of high port count port blades with port indexes greater than 256.

2

The unique area assignments are based on the port index. This mode does not allow FICON users to make use of ports with an index greater than 256 (high ports of a high port count blade), but this mode is compatible with domain-index zoning. This parameter is configurable only if the switch is Virtual Fabric-aware and Virtual Fabric is enabled on the switch.

R_A_TOV

The resource allocation time out value specified in milliseconds. This variable works with the variable E_D_TOV to determine switch actions when presented with an error condition.

Allocated circuit resources with detected errors are not released until the time value has expired. If the condition is resolved prior to the time out, the internal time-out clock resets and waits for the next error condition.

E_D_TOV

Error detect time out value specified in milliseconds. This timer is used to flag a potential error condition when an expected response is not received within the set time limit. If the time for an expected response exceeds the set value, then an error condition occurs.

WAN_TOV

Wide area network time out value specified in milliseconds. This timer is the maximum frame time out value for a WAN, if any, interconnecting the Fibre Channel islands.

MAX_HOPS

Maximum hops is an integer that denotes the upper limit on the number of hops a frame might have to traverse to reach any destination port from any source port across the fabric.

Note that the R_A_TOV, E_D_TOV, WAN_TOV, and MAX_HOPS configuration parameters are interrelated. Assigning a specific value to one or more of these parameters can change the range of allowed values that can be assigned to the other parameters. As a result, you may not be able to set all the values within the range displayed against each parameter. To reduce problems, the configuration utility validates the modified parameter values and prompts you to re-enter some values, if the validation check fails.

Data Field Size

The data field size specifies the largest possible value, in bytes, for the size of a type 1 (data) frame. The switch advertises this value to other switches in the fabric during construction of the fabric as well as to other devices when they connect to the fabric. Setting this parameter to a value smaller than 2112 might result in decreased performance.

Sequence-Level Switching

When sequence-level switching is set to 1, frames of the same sequence from a particular source are transmitted as a group. When this feature is set to 0, frames are transmitted interleaved among multiple sequences.

Under normal conditions, sequence-level switching should be disabled for better performance. However, some host adapters have performance issues when receiving interleaved frames from multiple sequences. When there are such devices attached to the fabric, sequence-level switching should be enabled.

Disable Device Probing

When disable device probing is set to 1, devices that do not register with the Name Server will not be present in the Name Server data base. Set this mode only if the switch's N_Port discovery process (PLOGI, PRLI, INQUIRY) causes an attached device to fail.

Suppress Class F Traffic

By default, the switch can send Class F frames. When this option is turned on, Class F traffic is converted to Class 2 traffic before being transmitted.

Per-frame Route Priority

In addition to the eight virtual channels used in frame routing priority, support is also available for per-frame-based prioritization when this value is set. When Per-frame Route Priority is set to 1, the virtual channel ID is used in conjunction with a frame header to form the final virtual channel ID.

Long Distance Fabric

When this mode is set to 1, ISLs in a fabric can be up to 100 km long. The exact distance level is determined by the per-port configuration on the E_Ports of each ISL. Both E_Ports in an ISL must be configured to run the same long-distance level; otherwise, the fabric will be segmented.

An Extended Fabrics license is required to set this mode.

BB Credit

The buffer-to-buffer (BB) credit represents the number of buffers available to attached devices for frame receipt. The range of allowed values varies depending on other system settings (see Unicast-only Operation). Refer to the *Fabric OS Administrator's Guide* for more information on platform-specific BB Credit limitations.

Insistent Domain ID Mode

When this mode is set, the switch attempts to acquire from the fabric the domain number programmed in its "Switch Fabric Settings." If the operation fails, the switch will segment from the fabric. You must disable the switch before configuring this parameter.

Disable Default PortName

When this mode is set, the switch does not generate a default port name.

Display FDMI Host Name

When this mode is set, the switch displays the FDMI Host name. By default this feature is in OFF state.

Dynamic Portname

When enabled, this feature dynamically assigns port name with various fields such as switch name, port type, port index, and alias name. Dynamic Portname and Display FDMI Host Name features are mutually exclusive.

Remote Fosexec feature

The remote fosexec configuration is bi-directional. The configuration is checked when sending any fosexec request to a remote switch and also when receiving the request from a remote switch. Thus, both the sending and receiving switch must be configured with fosexec ON. By default remote fosexec feature is in OFF state.

High Integrity Fabric Mode

Disables or enables the system to check for FMS mode and all existing HIF parameters present in the switch. By default this feature is in OFF state.

Edge hold time

Configures the maximum time a frame can wait after it is received on the ingress port and before it is delivered to the egress port. If the frame waits in the egress buffer for more than the configured hold time, the switch drops the frame, replenishes sender's credit, and increments the counters er_tx_c3_timeout and er_rx_c3_timeout on the TX and RX ports respectively. The frame-timeout indicates a slow draining or a congestion or bottleneck in the fabric. Decreasing hold time on the edge switches may reduce frame drop counts in the core switches. This parameter is stored persistently in the configuration file. You can configure edge hold time on both default and logical switch.

The edge hold time configuration is a chip-based value and the configuration varies between 8 Gbps-capable, 16 Gbps-capable, and 32 Gbps-capable platforms. 8 Gbps-capable platform provides one register to store the edge hold time value and any change in this value will affect all ports in the chip. Therefore, the edge hold time configuration will not be changed for 8 Gbps-capable platform ports in a logical switch. The 8 Gbps-capable platform ports in a logical switch take the default edge hold time value (220 milliseconds) or the value configured in the default switch. Both 16 Gbps and 32 Gbps-capable platforms provide four registers to store the edge hold time configuration and therefore each port in a chip can be configured one of the four values.

The edge hold time parameter is enabled by default with a value of 220 milliseconds. The following are the three predefined edge hold time values:

80

Low edge hold time (in milliseconds).

220

Medium edge hold time (in milliseconds). This is the default value.

500

High edge hold time (in milliseconds).

User Defined

A user-defined value may be defined in the range of 80 - 500 milliseconds in increments of one. This value is only applicable to the default switch.

In a logical switch, the edge hold time configuration is updated only for 16 Gbps-capable ports. If 8 Gbps-capable ports are present in the logical switch, the SWCH-1025 RASLog message is triggered to indicate that 8 Gbps-capable ports are present and the edge hold time configuration will not change for these ports.

F-Port Device Update Mode

When enabled, this feature provides the ability to update NS entries when the base devices perform logout. This option can only be changed when a switch is disabled. In Fabric OS prior to 7.3.0, the details of devices logged in with FLOGI (base devices) with AL_PA 0x40 and 0xCO are not propagated to NS or FCP unless the devices are registered with NS. Since they cannot be probed, devices with these AL_PAs must initiate NS registration. If they do register with NS and then later logout, there can be multiple LUNs losing access during login time because the logged out base device entry is not removed from NS. Furthermore, the logged out base device cannot login to the fabric again. Enabling this update mode ensures NS and login database consistency.

Virtual Channel Settings

VC Priority specifies the class of frame traffic given priority for a virtual channel. The switch enables fine-tuning for a specific application by configuring the parameters for eight virtual channels. The first two virtual channels are reserved for switch internal functions and are not available for modification.

The default virtual channel settings have already been optimized for switch performance. Changing the default values can improve switch performance but can also degrade performance. Do not change these settings without fully understanding the effects of the changes.

The values for virtual channel settings are as follows:

| Field | Default | Range |
|---------------|---------|--------|
| VC Priority 2 | 2 | 2 to 3 |
| VC Priority 3 | 2 | 2 to 3 |
| VC Priority 4 | 2 | 2 to 3 |
| VC Priority 5 | 2 | 2 to 3 |
| VC Priority 6 | 3 | 2 to 3 |
| VC Priority 7 | 3 | 2 to 3 |

F_Port Login Parameters

Specifies the F_Port login parameters. The following F_Port login settings are configurable. Unless there are issues with F_Port staging, do not change default values. (* = multiplication symbol)

| Field | Type | Default | Range |
|---------------------------------------|--------|---|-----------------------|
| Maximum logins per switch | Number | For directors: 16 * the max number of physical ports | 1 to 126 * the max |
| Logins per second | Number | 0 | 0 to 100 |
| Login stage interval (milli-seconds) | Number | 0 | 0 to 10000 |
| Stage FDISC logins with busy rejects: | Number | 0 | 1 to 255 |
| Enforce FLOGI/FDISC login: | Number | 0 | 0 to 1 |
| MAX num. of FLOGIs allowed | Number | 100 | 0 to 100 |

Maximum logins per switch

Sets a switch-wide limit on allowed logins.

The following three parameters are related to staged F_Port logins by FLOGI requests and virtual device logins by FDISC(SID==0) requests.

Logins per second

Specifies the number of logins the switch accepts per second in staged F_Port bring up.

Login stage interval

Specifies the stage interval in staged F_Port bring up.

Stage FDISC logins with busy reject:

This parameter, if nonzero, enables staging of FDISC logins by rejecting the FDISC requests with "logical busy", when the requests are more than the number of configured "logins per second". It also specifies the number of FDISC requests that will always be accepted first without reject.

Enforce FLOGI/FDISC login

Setting this flag allows a second F_Port login (FLOGI/FDISC login) in the event of two devices attempting to log in with the same PWWN. In default mode (zero) the first FLOGI/FDISC login takes precedence over the second. When the mode is set to 1, the second FLOGI/FDISC login takes precedence over the first. All modes are for NPIV and

non-NPIV F-Ports. When mode is set to 2, on FLOGI login the first FLOGI takes precedence. On FDISC login, the second FDISC takes precedence. For more information, refer to "Configure FLOGI-time handling of duplicate PWWNs" section in the *Fabric OS Administrator's Guide*. You must disable the switch to change this parameter.

MAX num. of FLOGIs allowed

Specifies the number of FLOGIs the port can accept per second. The default value is 100. When the FLOGI limit is reached, the port will be fenced.

Zoning Operation Parameters

The following zoning operation parameter can be modified.

Disable NodeName Zone Checking

Specify 1 to disable using node WWN when specifying nodes in the zone database. Specify 0 to enable using node WWN when specifying nodes in the zone data. The default value is 0. This value must be set to 1 for interpretability.

Local TI Filtering

Specify **on** to enforce TI zone rules on the local devices. The default value is **off**. You must first disable the switch using the **switchDisable** command before enabling this feature.

RSCN Transmission Mode

The RSCN transmission modes and values are as follows:

End-device RSCN Transmission Mode

Values are as follows:

| | |
|---|---------------------------------------|
| 0 | RSCN only contains single PID |
| 1 | RSCN contains multiple PIDs (Default) |
| 2 | Fabric addresses RSCN |

Domain RSCN to End-device

Values are as follows:

| | |
|---|--|
| 0 | Disabled. No domain RSCN is sent to the end-device for the switch IP address or name change. |
| 1 | Enabled. Domain RSCN is sent to the end-device for the switch IP address or name change. |

Arbitrated Loop Parameters

Arbitrated Loop Parameter and their values are as follows:

Send FAN frames?

Specifies that fabric address notification (FAN) frames be sent to public loop devices to notify them of their node ID and address. When set to 1, frames are sent; when set to 0, frames are not sent. The default value is 1.

Enable CLOSE on OPEN received?

If this is set, a CLS is returned immediately to an OPN if no buffers are available. This is required for TachLite. The valid range is 0 to 4. The default value is 0.

Always send RSCN?

Following the completion of loop initialization, a remote state change notification (RSCN) is issued when FL_Ports detect the presence of new devices or the absence of preexisting devices. When set to 1, an RSCN is issued upon completion of loop initialization, regardless of the presence or absence of new or preexisting devices. When set to 0, an RSCN is not sent (default).

System Services Settings

The values for the System Services settings are as follows:

Disable RLS probing

Enables or disables the read link status (RLS) probing. Performed by the FCP daemon, RLS probing reads the link error block from the device. This extended link services command is defined by the FC standards. Refer to the FC standards for information. RLS probing is enabled by default, meaning "Disable RLS probing" is "off". "on" disables RLS probing.

Portlog Events Enable/Disable Settings

These settings determine whether or not various types of port events are logged.

Each event type displayed on the screen is enabled by default ("on"). When disabled, this event is not logged by the port log.

Application Attributes

A number of application attributes are configurable. The following table specifies the attributes for applications that can be modified.

| Application | Field | Type | Default | Range |
|-------------|--------------------------------------|---------|-----------|--------------|
| ssl | Certificate File | string | [not-set] | varies |
| | CA Certificate File | string | [not-set] | varies |
| | Length of crypto key | number | [128] | 40,56,128 |
| rpcd | Secure RPCd Callback Secret | string | none | varies |
| Webtools | Basic User Enabled | boolean | no | no/yes |
| | Perform License Checking and Warning | boolean | yes | yes/no |
| | Allow Fabric Event Collection | boolean | yes | yes/no |
| | Login Session Timeout (in seconds) | number | 200 | 60 to 432000 |
| | | | | |

Notes The maximum per port login limit is no longer configurable with this command. Use the **portcfgNPiVPort --setloginlimit** command to configure this parameter on a per port basis.

The Telnet interface is no longer configurable with this command. Use the **ipFilter** command to enable or disable the Telnet interface.

The SNMP attributes are no longer configurable with this command. Use the `snmpConfig --set seclevel` command to configure SNMP attributes.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable XISL use on a logical switch with FID 20:

```
switch:admin> setcontext 20

switch:admin> switchdisable

switch:admin> configure
Configure...

Fabric parameters (yes, y, no, n): [no] y

    Domain: (1..239) [1]
    Enable a 256 Area Limit
      (0 = No,
       1 = Zero Based Area Assignment,
       2 = Port Based Area Assignment): (0..2) [0] 1
    WWN Based persistent PID (yes, y, no, n): [no]
Allow XISL Use (yes, y, no, n): [yes]yes
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000.. 5000) [2000]
WAN_TOV: (0..30000) [0]
MAX_HOPS: (7..19) [7]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]
Disable FID Check (yes, y, no, n): [no]
Insistent Domain ID Mode (yes, y, no, n): [no]
Disable Default PortName (yes, y, no, n): [no]
Display FDMI Host Name (yes, y, no, n): [no]
Edge Hold Time(Low(80ms), Medium(220ms), High(500ms)
  UserDefined(80-500ms): (80..500) [500]
Remote Fosexec feature: (on, off): [on]
High Integrity Fabric Mode (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
D-Port Parameters (yes, y, no, n): [no] yes
  Dynamic D-Port (on, off): [on]
  On Demand D-Port (on, off): [on]
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [1]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
```

configure

```
Portlog events enable (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]
switch:admin> switchenable
```

To enable Dynamic Area Mode on the default partition of the Brocade DCX:

```
switch:admin> switchdisable
switch:admin> configure
Configure...

Change fabric parameters? Y
  Domain: (1..239) [160]
  Enable 8 bit Dynamic Area Mode
    (0 = No,
     1 = Zero Based Area Assignment): (0..1) [0] 1
  R_A_TOV: (4000..120000) [10000]
  E_D_TOV: (1000..5000) [2000]
  WAN_TOV: (0..30000) [0]
  MAX_HOPS: (7..19) [7]
  Data field size: (256..2112) [2112]
  Sequence Level Switching: (0..1) [0]
  Disable Device Probing: (0..1) [0]
  Suppress Class F Traffic: (0..1) [0]
  Per-frame Route Priority: (0..1) [0]
  Long Distance Fabric: (0..1) [0]
  BB credit: (1..27) [16]
  Disable FID Check (yes, y, no, n): [no]
  Insistent Domain ID Mode (yes, y, no, n): [no]
  Disable Default PortName (yes, y, no, n): [no]
  Edge Hold Time(Low(80ms), Medium(220ms), High(500ms),
  UserDefined(80-500ms): (80..500) [500]
  Remote Fosexec feature: (on, off): [on]
  High Integrity Fabric Mode (yes, y, no, n): [no]
  Virtual Channel parameters (yes, y, no, n): [no]
  F-Port login parameters (yes, y, no, n): [no]
  D-Port Parameters (yes, y, no, n): [no] yes
    Dynamic D-Port (on, off): [on]
    On Demand D-Port (on, off): [on]
  RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [1]
  Zoning Operation parameters (yes, y, no, n): [no]
  RSCN Transmission Mode (yes, y, no, n): [no]
  Arbitrated Loop parameters (yes, y, no, n): [no]
  System services (yes, y, no, n): [no]
  Portlog events enable (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no]
  rpcd attributes (yes, y, no, n): [no]
  webtools attributes (yes, y, no, n): [no]
```

SEE ALSO [configDefault](#), [configShow](#), [configureChassis](#), [ipAddrSet](#), [portCfgLongDistance](#), [switchDisable](#), [switchEnable](#), [upTime](#)

configureChassis

Changes chassis-level system configuration settings.

Synopsis `configurechassis`

Description Use this command to modify chassis-level system configuration settings.

Configuration data is grouped into chassis information and switch information. Each configuration type is managed separately.

Use the **configure** command to modify switch configuration parameters. Use the **configureChassis** command to modify the following chassis configuration parameters:

- Cfgload attributes
- Custom Attributes
- System settings
- fos attributes

This command requires chassis permissions in both Virtual Fabric and non-Virtual Fabric environments.

The **configureChassis** command interactively presents a hierarchical menu. Each top-level heading and its associated subheadings consist of a text prompt, a selection of valid values, and a default value (in brackets).

The following keys control the execution of the command:

Return

When entered at a prompt with no preceding input, the command accepts the default value (if applicable) and moves to the next prompt.

Interrupt (Ctrl-C)

Aborts the command immediately and ignores all changes made.

End-of-file (Ctrl-D)

When entered at a prompt with no preceding input, terminates the command and saves changes made.

The following parameters can be modified with the **configureChassis** command:

Cfgload Attributes

Configures configuration upload and download parameters.

Enforce secure Config Upload/Download

Enables (yes) or disables (no) secure configuration upload or download. The default value is "no.

Add Suffix to the uploaded file name

Appends a suffix to the uploaded configuration file. The suffix includes the chassis name and a time stamp in the *yyymmdd_hhmmss* format. Refer to the example section for an illustration. This feature is disabled by default.

Do you want to enable auto firmwaresync

Enables the firmware from the active CP to synchronize automatically to the standby CP.

Custom Attributes

The following custom attributes can be modified.

config Index

OEM custom configuration. The range is 0 through 10000. The default is 0. This attribute is for internal use only.

System attributes

The following system-related parameters are configurable on a Virtual Fabric-aware switch.

system.blade.bladeFaultOnHwErrMsk

If this field is set to a value other than 0, then any nonfatal HW ASIC data parity error causes the problem blade to be powered off. The valid range is 0x0 to 0xffff. The default value is 0x0.

system.cpuLoad

Sets a threshold to define internally when the CPU is busy. The default threshold is 121, which represents a CPU instantaneous load average of 1.21 in the `top` command. The range is 10 to 121 (representing CPU load thresholds of .10 to 1.21).

system.i2cTurboCnfg

Configures the i2c driver that manages processing of the i2c interrupts. The I2C Turbo mode is enabled (Setting 1) by default on all 16G platforms running Fabric OS v7.0 or later. The mode is disabled by default for any switches capable of running Fabric OS v6.4x and Fabric OS v6.3x. An upgrade preserves the pre-upgrade configuration. The value shown in brackets is the current value. The following values are supported:

0

Disables the I2C Turbo mode. When the I2C Turbo mode is disabled, every byte received by an i2c interrupt is scheduled for processing by the given task or process. This permits other high priority processes to complete, thereby allowing for possible delay in the receipt of the I2C response message.

1

Enables the I2C Turbo mode for SFP transceivers only. This is the default setting. When the I2C Turbo mode is enabled, every I2C byte received from an SFP transceiver is processed by the interrupt handler itself, thus preventing any higher priority processes from preempting this processing.

2

For internal use only.

fos attributes

The following chassis-wide CS_CTL mappings can be modified:

CSCTL QoS Mode

Configures the CS_CTL to virtual channel (VC) mapping. After changing the CS_CTL QoS mode in a chassis, you must run the `slotPowerOff` and `slotPowerOn` commands for all the edge blades; In a fixed-port switch, you must reboot the switch. Because this mode change affects the persistent storage in the switch and chassis, rebooting is required for the new CS_CTL QoS mode to become effective.

- 0 Clears any previously configured CS_CTL to VC mapping and sets one-to-one mapping between a CS_CTL value and VC number for 8 Gbps-capable and 16 Gbps-capable ASIC. The CS_CTL values are divided into three groups: 1-8 (low priority), 9-16 (medium priority), and 17-24 (high priority). This is the default mode.
- 1 Sets the CS_CTL to VC mapping to support three CS_CTL values: 1 (low priority), 2 (medium priority), and 3 (high priority). Each CS_CTL value can map to more than one VC depending on the underlying support from ASIC. This is the auto mode. For more information about CS_CTL-based frame prioritization, refer to the *Fabric OS Administrator's Guide*.

Chassis SDDQ Limit

Specifies the number of user ports that can be quarantined in the chassis or unit. Valid values are 0 through 32. The default value is 10. The slow drain device quarantine (SDDQ) feature automatically isolates the slow-drain flows to a low priority VC from the existing VC (medium or high) thus freeing up the resources for the regular flows in the existing VC.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To add a suffix to the uploaded file name:

```
switch:admin> configurechassis
Configure...
cfgload attributes (yes, y, no, n): [no] y
    Enforce secure config Upload/Download (yes, y, no, n): [no]
    Add Suffix to the uploaded file name (yes, y, no, n): [no] y
    Do you want to enable auto firmwaresync (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
system attributes (yes, y, no, n): [no]
fos attributes (yes, y, no, n): [no]
```

To change the CS_CTL-based frame prioritization to default mode (one-to-one mapping):

```
switch:admin> configurechassis
Configure...
cfgload attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
system attributes (yes, y, no, n): [no]
fos attributes (yes, y, no, n): [no] y
    CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [1] 0
    Chassis SDDQ Limit: (0..32) [10]
```

configureChassis

To change the slow drain device quarantine limit:

```
switch:admin> configurechassis
Configure...

cfgload attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
system attributes (yes, y, no, n): [no]
fos attributes (yes, y, no, n): [no] y
  CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [1]
  Chassis SDDQ Limit: (0..32) [10] 20
```

SEE ALSO [configDefault](#), [configShow](#), [chassisEnable](#), [chassisDisable](#), [configure](#), [ipAddrSet](#), [portCfgLongDistance](#), [switchDisable](#), [switchEnable](#), [upTime](#)

creditRecovMode

Enables credit loss recovery on back-end ports of 8Gbps-capable, 16Gbps-capable, 32 Gbps-capable blades, and time-out based credit loss detection of 8 Gbps-capable front-end ports.

Synopsis

```
creditrecovmode --cfg [off | onLrOnly | onLrThresh]
  [-lrthreshold threshold]

creditrecovmode --fault [edgeblade | coreblade | edgecoreblade]

creditrecovmode --check [slot]/blade_port,VC

creditrecovmode --fe_crdloss [off | on]

creditrecovmode --be_crdloss [off | on]

creditrecovmode --be_losync [off | on]

creditrecovmode --linkreset [slot]/blade_port

creditrecovmode --show

creditrecovmode --help
```

Description Use this command to enable or disable credit recovery of backend ports and to display the configuration. When this feature is enabled, credit is recovered on backend ports (ports connected to the core blade or core blade backend ports) when credit loss has been detected on these ports. If complete loss of credit on a 8 Gbps-capable backend port causes frame timeouts, a link reset will be performed on that port regardless of the configured setting, even if that setting is **--cfg off**. When used with the **--cfg onLrOnly** option, the recovery mechanism takes the following escalating actions:

- When it detects credit loss, it performs a link reset and logs a RASlog message (RAS Cx-1014).
- If the link reset fails to recover the port, the port reinitializes. A RASlog message is generated (RAS Cx-1015). Note that the port reinitialization does not fault the blade.
- If the port fails to reinitialize, the port is faulted. A RASlog message (RAS Cx-1016) is generated.
- If a port is faulted and there are no more online backend ports in the trunk, the core blade is faulted. (Note that the port blade will always be faulted.) A RASlog message is generated (RAS Cx-1017).

When used with the **--cfg onLrThresh** option, recovery is attempted through repeated link resets and a count of the link resets is kept. If the threshold of more than the configured threshold value (using the **-lrthreshold** option) per hour is reached, the blade is faulted (RAS Cx-1018). Note that regardless of whether the link reset occurs on the port blade or on the core blade, the port blade is always faulted.

The **onLrOnly** and **onLrThresh** options activate link reset for both credit loss and loss of synchronization. The **be_crdloss** option activates link reset for credit loss only. The **be_losync** option activates link reset for loss of synchronization only.

If you suspect complete credit loss on a particular virtual channel (VC) on a particular backend port, use the **--check** option to examine that particular backend port and VC for credit loss. If the command detects complete credit loss, it reports the information. If, in addition, you have enabled link resets on backend ports, this command will perform a link reset on the link in an attempt to recover from the problem. You must explicitly initiate this check and it is a one-time operation. In other words, this command does not continuously monitor for credit loss in the background. Detection of credit loss takes 2 to 7 seconds, after which the results of the operation are displayed. A link reset also generates a RASlog message.

For more information on the RASlog messages, refer to the *Fabric OS Message Reference*.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported only on backend ports of 8 Gbps-capable, 16 Gbps-capable, and 32 Gbps-capable blades.

The **--check** option is supported only on modular switches, and only on links between 8 Gbps-capable backend ports or the 8 Gbps-capable to 16 Gbps-capable backend ports. AP blades are not supported.

The **--fe_crdloss** option is supported only on 8 Gbps-capable front-end interswitch link (ISL) ports.

Operands This command has the following operands:

--cfg

Configures credit recovery for backend ports. Use one of the following required recovery options to configure credit recovery:

onLrOnly

Enables the backend port recovery feature in link reset mode.

onLrThresh

Enables the backend port recovery feature in link reset threshold mode.

off

Disables the backend port credit recovery feature.

-lrthreshold *threshold*

Specifies the link reset threshold value. The default value is 2. Note that the threshold value set using this option will apply only when the recovery mode is **onLrThresh**.

--fault

Configures the blade fault option. Valid values are as follows:

edgeblade

The edge blade alone is faulted when the link reset threshold is reached.

coreblade

The core blade alone is faulted when the link reset threshold is reached.

edgecoreblade

The edge blade is faulted the first two times the link reset threshold value is reached. If the threshold value is reached the third time, the core blade is faulted.

| | |
|---|--|
| --fe_crdloss | Configures time-out based credit loss detection of 8 Gbps-capable front-end ISL links. This feature is disabled by default. |
| off | Disables credit loss detection on 8 Gbps-capable front-end ports. |
| on | Enables credit loss detection on 8 Gbps-capable front-end ports. |
| --be_crdloss | Activates link reset when credit loss occurs on backend ports. This feature is disabled by default. |
| off | Disables link reset for credit loss on backend ports. |
| on | Enables link reset for credit loss on backend ports. |
| --be_losync | Activates link reset when loss of synchronization occurs on backend ports. This feature is disabled by default. |
| off | Disables link reset for loss of synchronization on backend ports. |
| on | Enables link reset for loss of synchronization on backend ports. |
| -check | Performs the on-demand detection of credit loss on a given backend port and VC. This operand is optional and exclusive; when you specify this option you cannot use the --cfg options at the same time. You must specify a port and a VC, separated by a comma. A space before the comma is not permitted. |
| <i>slot/blade_port</i> | Specifies the backend port that is to be examined for credit loss. The port number must be the blade port number, because this is a backend port. The blade port number can be located in the "Bpt" column output of the bladePortMap command. Note that the bladeportmap command requires root permissions. |
| VC | Specifies the Virtual Channel number. The valid range is from 1 through 31. VC 0 is invalid. |
| --linkreset <i>slot/blade_port</i> | Performs a link reset on the specified front-end or back-end blade port. The blade port number can be located in the "Bpt" column output of the bladePortMap command. The bladeportmap command requires root permissions. |
| --show | Displays the backend port credit recovery configuration as enabled or disabled. In addition, the output indicates whether link reset mode or link reset threshold mode is configured. |
| --help | Displays the command usage. |

Examples

To enable backend port credit loss recovery with the link reset only option and to display the configuration:

```
switch:admin> creditrecovmode --cfg onLrOnly
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrOnly
C2 FE Complete Credit Loss Detection is Enabled
```

To enable backend port credit loss recovery with the link reset threshold option and to display the configuration:

```
switch:admin> creditrecovmode --cfg onLrThresh
switch:admin> creditrecovmode --show
Internal port credit loss recovery is Enabled with LrThresh
C2 FE Complete Credit Loss Detection is Enabled
```

To disable backend port credit loss recovery and to display the configuration:

```
switch:admin> creditrecovmode --cfg off
switch:admin> creditrecovmode --show
Internal port credit loss recovery is Disabled
C2 FE Complete Credit Loss Detection is Enabled
```

To configure a fault option and to display the configuration:

```
switch:admin> creditrecovmode --fault edgecoreblade
switch:admin> creditrecovmode --show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : EDGECOREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

To disable credit loss detection on 8 Gbps-capable front-end ports:

```
switch:admin> creditrecovmode --fe_crdloss off
switch:admin> creditrecovmode --show
Internal port credit loss recovery is Disabled
C2 FE Complete Credit Loss Detection is Disabled
```

To examine a backend port for credit loss:

```
switch:admin> creditrecovmode --cfg off
switch:admin> creditrecovmode --check 2/32,1
Started Credit loss Detection on slot 2 port 32 VC 1.
Please wait....
Detected credit loss.
Link Reset performed.
```

To activate credit loss and loss of sync detection with link reset and set link reset threshold value to 10:

```
switch:admin> creditrecovmode --cfg onLrThresh --lreshold 10
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrThresh
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 10
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

To disable link reset for credit loss alone:

```
switch:admin> creditrecovmode --be_crdloss off
switch:admin> creditrecovmode --show
Internal port credit recovery is Disabled
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

To enable link reset for loss of sync alone:

```
switch:admin> creditrecovmode --be_losync on
switch:admin> creditrecovmode --show
Internal port credit recovery is Enabled with LrThresh
Back end port Loss of Sync's Link Reset is Enabled with LrThresh
LR threshold (currently activated): 2
Fault Option : COREBLADE
C2 FE Complete Credit Loss Detection is Disabled
```

SEE ALSO None

dataTypeShow

Displays sample data stream types used in some diagnostic commands.

Synopsis `datatypeshow [-seed value]`

Description Use this command to display sample data stream types used in diagnostic commands. There are 25 different sample data types. The command displays an example of each data stream.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

-seed *value* Specify the data pattern seed value. This operand is optional. The default value is 0.

Examples To display sample data streams you can use with diagnostics:

```
switch:admin> datatypeshow

Pattern  type      example
-----
BYTE_FILL  1      15 15 15 15 15 15 15 15 15 15 15 15 15 15 15 15
WORD_FILL  2      0015 0015 0015 0015 0015 0015 0015 0015 0015
QUAD_FILL  3      00000015 00000015 00000015 00000015
BYTE_NOT   4      15 ea 15 ea 15 ea 15 ea 15 ea 15 ea 15 ea 15 ea
WORD_NOT   5      0015 ffea 0015 ffea 0015 ffea 0015 ffea
QUAD_NOT   6      00000015 fffffffea 00000015 fffffffea
BYTE_RAMP  7      15 16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24
WORD_RAMP  8      0015 0016 0017 0018 0019 001a 001b 001c
QUAD_RAMP  9      00000015 00000016 00000017 00000018
BYTE_LFSR  10     1c 38 71 e3 c6 8d 1a 34 68 d0 a0 41 82 04 09 12
RANDOM      11     14 5b 52 66 77 c5 bc 76 5a 6e d3 e2 de 3b 6b 40
CRPAT      12     bc bc 23 47 6b 8f b3 d7 fb 14 36 59 bc bc 23 47
CSPAT      13     7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f
CHALF_SQ   14     b5 b5 b5 b5 b5 b5 b5 b5 b5 b5 b5 b5 b5 b5 b5
CQTR_SQ    15     d9 26 d9 26 d9 26 d9 26 d9 26 d9 26 d9 26 d9
RDRAM_PAT  16     ff 00 ff 00 ff 00 ff 00 ff 00 ff 00 ff 00 ff 00
jCRPAT     17     be d7 23 47 6b 8f b3 14 5e fb 35 59 be d7 23 47
jCJTPAT    18     7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e
jCSPAT     19     7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f 7f
PRED_RAND  20     00000015 11111126 22222237 33333348
SMI_TEST   21     00000015 fffffffea 00000015 fffffffea
CJPAT      22     7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e 7e
QUAD_NOTP  23     00000015 ffffffffa 00000015 ffffffffa
JSPAT      24     7e 7e 7e 7e 7e 7e 7e 9e f8 50 81 e7 50 aa c9
JTSPAT     25     7e 7e 7e 7e 7e 7e 7e 9e f8 50 81 e7 50 aa c9
```

SEE ALSO None

date

Displays or sets the switch date and time.

Synopsis `date ["newdate"]`

Description Use this command to display or set the date and time. All switches maintain current date and time in flash memory. Date and time are used for logging events. Normal switch operation does not depend on the date and time; a switch with incorrect date values continues to function properly.

This command sets a common date and time for the entire fabric. A change in date or time to one switch is forwarded to the principal switch and distributed to the fabric. It may take up to 64 seconds for the switches in the fabric to be synchronized. However, if an FCS policy is enabled, this command can be executed only on the Primary FCS switch, and only the primary FCS switch can distribute the time stamp to all other switches in the fabric.

If Virtual Fabrics are enabled, the date is set for the entire chassis, including all logical switches.

The date specified is always the local switch time, taking into account daylight saving time and the time zone setup of the switch. Each switch takes care of converting the GMT time distributed fabric-wide to its local time. Refer to **tsTimeZone** for more information on time zone support.

If the switch is operating in FICON Management Server mode (**fmsMode**), setting the date is subject to the director clock alert mode (DCAM). If DCAM is 1, the operator issues a warning that the switch date is about to change. The operator then prompts to confirm the change with a yes or no response.

Notes This command becomes read-only if external NTP synchronization is enabled. For more information, refer to **tsClockServer**.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

| | |
|-------------------------------|--|
| <code>"<i>newdate</i>"</code> | Specify the new date and time enclosed in double quotation marks. this operand is optional; if omitted, the current date and time is displayed. Date and time are specified as a string in the <i>mmddhhmmyy</i> format. |
| <code><i>mm</i></code> | Specifies the month. Valid values are 01 to 12. |
| <code><i>dd</i></code> | Specifies the date. Valid values are 01 to 31. |
| <code><i>hh</i></code> | Specifies the hour. Valid values are 00 to 23. |
| <code><i>mm</i></code> | Specifies the minutes. Valid values are 00 to 59. |

date

yy

Specifies the year, valid values are 00 to 37 and 70 to 99. Year values from 70 to 99 are interpreted as 1970 to 1999; year values from 00 to 37 are interpreted as 2000 to 2037.

Examples To display the current date and time and then modify it:

```
switch:admin> date
Tue Oct 22 14:05:10 UTC 2013
switch:admin> date "1022140613"
Tue Oct 22 14:06:00 UTC 2013
```

SEE ALSO [errShow](#), [ficonCupSet](#), [ficonCupShow](#), [portLogShow](#), [tsClockServer](#), [tsTimeZone](#), [upTime](#)

dbgShow

Displays current values of debug and verbosity levels of the specified module.

Synopsis `dbgshow [module_name]`

Description Use this command to display the current values of debug and verbosity levels of the specified module. If no module name is specified, this command displays a listing of all modules along with debug and verbosity levels.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

module_name Specifies the name of the module for which you want to view the debug and verbosity levels. Module names are case-sensitive. This operand is optional.

Examples To display information about a specific module named NS:

```
switch:admin> dbgshow NS
Module NS,      debug level = 1, verbose level = 1
```

SEE ALSO [setDbg](#)

defZone

Sets or displays the default zone access mode.

Synopsis `defzone [--noaccess | --allaccess | --show]`

Description Use this command to display or set the Default Zone access mode. Setting the Default Zone mode initializes a zoning transaction (if one is not already in progress), and create reserved zoning objects.

A default zone controls device access when zoning is not enabled. When a user-specified zoning configuration is not enabled, Default Zone is in effect, allowing access to all devices. When a user-specified zone configuration is enabled, it overrides the Default Zone access mode.

Notes This command must be issued from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Zone object names beginning with the `d_default_` prefix are reserved for default zoning use. Editing of these objects is not permitted. Therefore, `cfgShow` does not display the names of these objects.

If `d_default_Cfg` is the effective zone configuration, both `cfgShow` and `cfgActvShow` do not display `d_default_Cfg` as the effective zone configuration.

Operands This command has the following operands:

--noaccess

Sets the default zone access mode to No Access, initializes a zoning transaction (if one is not already in progress), and creates the reserved zoning objects equivalent to the following zoning commands:

- `cfgCreate "d_default_Cfg","d_default_Zone"`
- `zoneCreate "d_default_Zone","00:00:00:00:00:00:00:01"`

A `cfgSave`, `cfgEnable`, or `cfgDisable` command must be issued after issuing this command to commit the changes and distribute them to the fabric; for example:

- `defzone --noaccess`
- `cfgsave`

An audit log record is generated for each execution of this command.

When No Access default zone is activated, the following conditions apply:

- If the current effective zone configuration is disabled with the **cfgDisable** command, the local switch converts this command to the equivalent of **cfgEnable "d_default_Cfg"**.
- If zoning receives a **cfgDisable** command from a remote switch that does not support default zoning, zoning rejects the **cfgDisable** command in the second phase of RCS because the remote switch does not convert the **cfgDisable** command to **cfgEnable "d_default_Cfg"**.

--allaccess

Sets the default zone access mode to All Access, initiates a zoning transaction (if one is not already in progress), and deletes the reserved zoning objects by performing the equivalent to the following zoning commands:

- `cfgDelete "d_default_Cfg"`
- `zoneDelete "d_default_Zone"`

A **cfgSave**, **cfgEnable**, or **cfgDisable** command must be performed subsequent to the use of this command to commit the changes and distribute them to the fabric. If a **cfgSave** is performed and the fabric is already in the No Access default zone state, a **cfgDisable** is sent to the fabric. For example:

- `defzone --allaccess`
- `cfgsave`

An audit log record is generated for each use of this command.

--show

Displays the current state of the default zone access mode.

Examples

To create a default zone configuration:

```
primaryfcs:admin> cfgactvshow
Effective configuration:
No Effective configuration

primaryfcs:admin> defzone --noaccess

primaryfcs:admin> cfgsave

primaryfcs:admin> defzone --show
Default Zone Access Mode
committed - No Access
transaction - No Transaction

primaryfcs:admin> cfgactvshow
Effective configuration:
No Effective configuration: (No Access)
```

SEE ALSO None

diagClearError

Clears the diagnostics failure status.

Synopsis `diagclearerror [--slot slot] -all`

Description Use this command to clear the diagnostics failure status. When used without operands, this command clears all port failure flags.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

- slot *slot*** Specify the slot on which to clear the diagnostics failure status. The default is set to 0 and designed to operate on fixed-port-count products.
- all** If specified, all blades clear.

Examples To clear the diag software flag:

```
switch:admin> diagclearerror --slot 8
                ERROR: DIAG CLEARERR
Diagnostics Errors Cleared, port: 8/31
Err# 0120041 081F
```

SEE ALSO None

diagDisablePost

Disables the power-on self-test (POST).

Synopsis `diagdisablepost`

Description Use this command to disable POST. A reboot is not required for this command to take effect. Use the **diagPost** command to display the current POST status, and use **diagEnablePost** to enable POST.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the POST:

```
switch:admin> diagdisablepost  
Config update Succeeded  
Diagnostic POST is now disabled.
```

SEE ALSO [diagEnablePost](#), [diagPost](#)

diagEnablePost

Enables the power-on self-test (POST).

Synopsis `diagenablepost`

Description Use this command to enable POST. A reboot is not required for this command to take effect. POST includes two phases: POST Phase I mainly tests hardware and POST Phase II tests system functionality.

Use the **diagPost** command to display the current POST status, and use **diagDisablePost** to disable POST.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable POST:

```
switch:admin> diagenablepost
Config update Succeeded
Diagnostic POST is now enabled.
```

SEE ALSO [diagDisablePost](#), [diagPost](#)

diagHelp

Displays diagnostic command information.

Synopsis `diaghelp`

Description Use this command to display a short description of diagnostic commands.

Use default operands when running diagnostics commands. Non-default settings require detailed knowledge of the underlying hardware and are intended for support personnel only. Contact support if you want to use these operands.

Notes The **diagHelp** command displays diagnostic commands that may not be available. Execute **help command** to verify availability.

Operands None

Examples To display diagnostic command information:

```
switch:admin> diaghelp
bportloopbacktest      Functional test of port via
                        blade processor path.
bpturboramtest         MBIST test for AP blade BP ASICs
ceepportloopbacktest  Functional test of port N->N path.
ceeturboramtest        MBIST test for ASICs
(output truncated)
```

SEE ALSO None

diagPost

diagPost

Displays the diagnostic power-on self-test (POST) configuration.

Synopsis `diagpost`

Description Use this command to display the current POST configuration. Use **diagEnablePost** or **diagDisablePost** to modify the POST configuration.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the current POST configuration:

```
switch:admin> diagpost  
Diagnostic POST is currently disabled.
```

SEE ALSO None

diagRetry

Sets or displays diagnostic retry mode.

Synopsis `diagretry [mode | -show]`

Description Use this command to enable retry mode if the mode value is nonzero and to disable the retry mode if the mode value is 0. The mode value is saved in nonvolatile memory until you change the mode. Changes made by this command do not require a reboot to take effect.

Retry mode modifies the behavior of the diagnostic test methods, power-on self-test (POST), and burn-in scripts. The exact behavior depends on the tests and scripts that are run.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

mode

Specify 1 to enable retry mode. Specify 0 to disable the feature.

-show

Displays the current mode setting. The **--show** option is equivalent to using the command without operands.

Examples To view and modify the current retry mode value:

```
switch:admin> diagretry -show
Diagnostic Retry Mode is currently enabled.
```

```
switch:admin> diagretry 0
Config update Succeeded
```

```
Diagnostic Retry Mode is now disabled.
```

```
switch:admin> diagretry 1
Config update Succeeded
Diagnostic Retry Mode is now enabled.
```

SEE ALSO None

diagShow

Displays diagnostics status.

Synopsis

```
diagshow
  [--slot number]
  [-post]
  [-port number]
  [-use_bports value]
```

Description Use this command to display the diagnostics status for the specified list of ports.

Notes You cannot interrupt the test by pressing the return key (<cr>).

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

--slot *number* Specify which slot to operate on. If this option is not specified, the default slot 0 is used. The default slot is designed to operate on fixed-port-count products. By default, this command displays information for all user ports in the system. Default slot 0 is not valid in a chassis system.

-post Displays the status of the last run power-on self-test (POST) on the specified slot. This operand is supported only on 32 Gbps-capable platforms.

-port *number* Displays status of the specified user slot. This is the default behavior if the **-post** option is not specified.

-use_bports *value* If this value is nonzero, this command displays the diagnostics status for the blade ports specified in **-use_bports**; otherwise, the command displays information for the user ports specified in **-uports**. The default value is 0.

Examples To display diagnostic status on a switch blade:

```
switch:admin> diagshow
Slot#           Switch Type           ID   Status  Rev#   Blade   Post Status
0               Fixed Configuration  171  ENABLED  2.0   FX8-24  POST PASSED
(output truncated)
```

SEE ALSO [itemList](#)

diagStatus

Displays currently running diagnostic tests.

Synopsis `diagstatus [slotnumber]`

Description Use this command to display currently running diagnostic test names.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

slotnumber

Specifies the slot number to display diagnostic test names. If omitted, information for all blades in the system is displayed.

Examples To display currently running diagnostic tests:

```
switch:user> diagstatus
Diagnostic status for slot: 1.
Diag executing "NONE"
Diagnostic status for slot: 2.
Diag executing "NONE"

Diagnostic status for slot: 3.
Diag executing "NONE"
Diagnostic status for slot: 4.
Diag executing "NONE"
(output truncated)
```

SEE ALSO None

distribute

distribute

Distributes data to switches in a fabric.

Synopsis `distribute -p policy_list -d switch_list`

Description Use this command to distribute data to a specified list of switches in the fabric. The distributed data must be from the list of currently supported policy sets:

| | |
|------|---|
| SCC | Switch Connection Control Policy |
| DCC | Device Connection Control Policy |
| PWD | Password Database and Password Configuration Policy |
| AUTH | E_Port and F_Port Authentication Policy |
| FCS | Fabric Configuration Server Policy |

Each supported database has a switch-local configuration parameter that controls whether the database can be distributed and accepts distributions. Use the **fddCfg** command to view and modify these parameters.

Notes IP Filter policies cannot be distributed with the **distribute** command. Use the **chassisDistribute** command.

The password database received from a switch running pre-v6.3.0 firmware will be rejected by a Virtual Fabric-enabled v6.3.0 chassis if it has more than one logical switch.

If FCS policy is enabled, only primary FCS switch can distribute the data.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-p *policy_list* Specify the list of policy sets, also called security databases, to be distributed. *policy_list* is a semicolon-separated list. Valid values include SCC, DCC, PWD, AUTH, and FCS.

-d *switch_list* Specify the list of switches that should receive the data distribution. The *switch_list* is a semicolon-separated list of one of the following:

- Switch domain IDs
- Switch names
- Switch WWNs

A wildcard (*) may be specified to include all switches in the fabric that support the **distribute** feature.

Examples To distribute the Switch Connection Control Policy and Device Connection Control Policy to domains 3 and 5 in the fabric:

```
switch:admin> distribute -p "SCC;DCC" -d "3;5"
```

To distribute the Switch Connection Control Policy, Fabric Configuration Server Policy, and Password database to all domains in the fabric that support the **distribute** feature:

```
switch:admin> distribute -p "SCC;FCS;PWD" -d ""  
Wildcard domains are:  
1 3 5
```

To distribute the FCS policy, and the Password database to all domains in the fabric that support the **distribute** feature:

```
switch:admin> distribute -p "FCS;PWD" -d ""
```

To distribute the AUTH and FCS policies to all switches in the fabric:

```
switch:admin> distribute -p "AUTH;FCS" -d ""
```

To distribute the AUTH and SCC policies to domains 1 and 3 in the fabric:

```
switch:admin> distribute -p "AUTH;SCC" -d "1;3"
```

SEE ALSO [fddCfg](#)

dlsReset

Disables Dynamic Load Sharing (DLS).

Synopsis `dlsreset`

Description Use this command to disable Dynamic Load Sharing.

If DLS is turned off, load sharing calculations are used only to place new routes. Once placed, existing routes are never moved from one output E_Port to another, unless the original output E_Port is no longer a recognized path to the remote domain. Optimal balance is rarely achieved with this setting. Refer to the **dlsSet** help page for a full description of load sharing options, including the Lossless feature and the E_Port Balance Priority feature.

The behavior of this command depends on the routing policies configured on the switch:

- If a port-based routing policy is in place, DLS is disabled by default, and **dlsReset** returns the DLS setting to default. This command resets the Lossless, and the E_Port Balance Priority features along with the DLS feature and returns a message stating that "DLS is not set".
- If an exchange-based routing policy is in place, DLS is always enabled. It cannot be disabled and the **dlsReset** command fails. The command generates a message stating that "DLS is enabled and cannot be changed with the current routing policy". If Lossless is enabled, the feature remains enabled until you disable it with the **dlsSet --disable -lossless** command. If E_Port Balance Priority is enabled, the feature remains enabled until you disable it with the **dlsset --disable -eportbal** command.
- If DLS is already disabled, the command output confirms the disabled status: "DLS is not set (unchanged)."
- If two-hop lossless DLS is enabled while DLS is disabled, the command output shows the message: "Two-hop lossless requires LosslessDLS be enabled prior to enabling."

Refer to **aptPolicy** for information on routing policies.

Notes The Lossless feature is not supported on GbE ports and FCoE ports. On unsupported platforms, this command exits with an appropriate message.

On platforms that do not support the E_Port Balance Priority feature, the command output indicates that the feature is not supported.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the Dynamic Load Sharing option on a switch with a port-based routing policy and DLS enabled:

```
switch:admin> dlsreset
DLS is not set
```

To execute **dlsReset** on a switch with an exchange-based routing policy and DLS enabled:

```
switch:admin> dlsreset
DLS is enabled and cannot be changed with the current routing policy
```

SEE ALSO [aptPolicy](#), [dlsSet](#), [dlsShow](#)

dlsSet

Enables Dynamic Load Sharing (DLS) without frame loss.

Synopsis

```
dlsset
dlsset --enable -lossless
dlsset --enable -eportbal
dlsset --enable -twohop
dlsset --disable -lossless
dlsset --disable -eportbal
dlsset --disable -twohop
dlsset --rebalance
dlsset --rebalance -all
dlsset --help
```

Description

Use this command to enable or disable lossless Dynamic Load Sharing (DLS) in the event of a fabric change, to configure DLS without frame loss, and to display the DLS configuration.

Use this command to enable or disable the E_Port Balance Priority feature. Enabling E_Port balancing causes the E_Port load to be even across all E_Ports to the same domain. E_Port balancing is disabled by default.

Dynamic load sharing optimizes the utilization of the interswitch links (ISLs) by rebalancing the paths going over the ISLs whenever there is a fabric event that may result in a sub-optimal utilization of the ISL. Dynamic rebalancing can be triggered by any one of the following events:

- A change in the fabric occurs.
- A local E_Port (including trunk ports) goes up or down.
- A local Fx_Port goes down.

When used without operands, this command enables Dynamic Load Sharing on a switch (legacy DLS behavior). Frames may be lost during reroute operations. If the switch has an exchanged-based routing policy, DLS is enabled by default and this command fails with the following message: "DLS is enabled and cannot be changed with the current routing policy."

During the load sharing recomputation, existing routes may be moved to maintain optimal load balance. Frame loss is unavoidable when a port goes down. To prevent frames from being lost during this operation, you can enable DLS without frame loss by issuing this command with the **--enable -lossless** option.

Dynamic load sharing without frame loss is supported in logical fabrics and is configured per logical switch. However, there is a potential impact on other logical switches because they share the same hardware. Chassis permissions are required to configure DLS in a logical fabric environment.

For example, assume a chassis is partitioned as follows: logical switch LS1 consists of ports 1/0-1/5, and logical switch LS2 consists of ports 1/6-1/10. Lossless is enabled on logical switch LS1. Because the ports 1/0-1/10 share the same chip, traffic in LS2 is affected whenever traffic for LS1 on ports 1/0-1/5 is rebalanced. The impact on LS2 depends on the configuration on LS2:

- If the Lossless feature is enabled on LS1, traffic pauses and resumes without frame loss on both switches at the same time.
- If the Lossless feature is disabled on LS1, traffic on LS2 is not affected.

Notes

For switches running Fabric OS v7.1.0 or later, you can enable Lossless (or enable Fabric Management Server mode) when XISL is enabled and enable XISL when Lossless or Fabric Management Server mode is enabled.

When you downgrade from Fabric OS v7.1.x to Fabric OS v7.0.x the following rules apply: If Lossless (or Fabric Management Server mode) and XISL use are not enabled at the same time, firmware download can be executed. If both of them are enabled, firmware download is rejected with an error message.

If the active control processor (CP) runs Fabric OS v7.1.x or later, and the standby CP runs Fabric OS v7.0.x and if both Lossless (or Fabric Management Server mode) and XISLs are enabled, High Availability synchronization will fail. Otherwise, standby CP will synchronize with the active CP. After the synchronization of the active and standby CP, enabling Lossless (or Fabric Management Server mode), when XISL is enabled, is rejected with an error message. Enabling XISL, when Lossless (or Fabric Management Server mode) is enabled, is also rejected with an error message.

Use the **configure** command to disable XISL use.

Lossless is supported with both exchange-based and port-based routing policies. Behavior depends on the kind of policy configured and concurrent IOD settings. Refer to the *Fabric OS Administrator's Guide* for information on how to optimize your configuration. Refer to the **aptPolicy** help page for more information on routing policies.

The Lossless feature is not supported on GbE ports and FCoE. On unsupported platforms, this command exits with an appropriate message.

When the E_Port Balance Priority feature status is changed using **--enable -eportbal** or **--disable -eportbal**, the status of the DLS and Lossless features remains unchanged. The E_Port Balance Priority feature is supported only on multi-ASIC platforms like Brocade 6520, DCX 8510-8, and DCX 8510-4.

Two-hop Lossless DLS is not supported over LISLs.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

--enable -lossless

Enables the Lossless feature. Frame loss is reduced while the path is rerouted. If DLS is set on the switch, this command adds the Lossless feature to the existing DLS legacy mode. If DLS is not set on the switch, this command enables both DLS and the Lossless feature. The Lossless feature is disabled by default.

--disable -lossless

Disables the previously enabled Lossless feature. Execution of this command is equivalent to the legacy **dlsSet** command. Dynamic load balancing is enforced but not without frame loss. DLS (legacy mode) continues to be enabled after Lossless is disabled. Use **dlsReset** to disable DLS completely.

- enable -eportbal**
Enables the E_Port Balance Priority feature. The E_Port load is rebalanced when topology changes occur, such as an E_Port going offline and then online. If DLS is disabled on the switch, it remains disabled with this command.
- disable -eportbal**
Disables the E_Port Balance Priority feature.
- enable -twohop**
Enables two-hop lossless DLS. When enabled, this enhancement coordinates route updates of the switches in the fabric such that the intermediate switches have performed their route rebalancing updates for the new path prior to an upstream switch attempting to use that new path. The Lossless feature must be enabled to perform route updates. Two-hop lossless DLS is disabled by default.
- disable -twohop**
Disables the previously enabled two-hop lossless DLS.
- rebalance**
Rebalances the E_Port load on a logical switch, without waiting for a topology change to occur.
- rebalance -all**
Rebalances the E_Port load on all logical switches, without waiting for a topology change to occur.
- help**
Displays the command usage.

Examples DLS configuration commands on a switch with an exchange-based policy:

```
switch:admin> aptpolicy
Current Policy: 3 0(ap)

3 0(ap) : Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
    0: AP Shared Link Policy
    1: AP Dedicated Link Policy

switch:admin> dlsshow

DLS is set by default with current routing policy

DLS is set with Lossless enabled

E-port Balance Priority is not set

switch:admin> dlsreset

DLS is enabled and cannot be changed with the current routing policy

switch:admin> dlsset

DLS is enabled and cannot be changed with the current routing policy
```

```
switch:admin> dlsset --enable -lossless
```

```
Lossless is set
```

```
switch:admin> dlsset --disable -lossless
```

```
Lossless is not set
```

```
switch:admin> dlsshow
```

```
DLS is set by default with current routing policy
```

```
E-port Balance Priority is not set
```

DLS configuration commands on a switch with a port-based policy:

```
switch:admin> dlsshow
```

```
DLS is set by default with current routing policy
```

```
DLS is set with Lossless enabled
```

```
E-port Balance Priority is not set
```

```
switch:admin> dlsreset
```

```
DLS is not set
```

```
switch:admin> dlsshow
```

```
DLS is not set
```

```
E-port Balance Priority is not set
```

```
switch:admin> dlsset --enable -lossless
```

```
DLS and Lossless are set
```

```
switch:admin> dlsshow
```

```
DLS is set with Lossless enabled
```

```
switch:admin> dlsset
```

```
Lossless is set (unchanged)
```

```
switch:admin> dlsset --disable -lossless
```

```
Lossless is not set
```

```
switch:admin> dlsshow
```

```
DLS is not set
```

```
E-port Balance Priority is not set
```

To attempt to enable Lossless while XISL use is enabled:

```
switch:admin> dlsset --enable -lossless
```

```
Lossless option cannot be enabled when XISL use is allowed.  
Please disable the switch with 'switchdisable' and run \  
'configure' to disallow XISL use before enabling Lossless.
```

To enable the E_Port Balance Priority:

```
switch:admin> dlsset --enable -eportbal  
E-port Balance Priority is set
```

To disable the E_Port Balance Priority:

```
switch:admin> dlsset --disable -eportbal  
E-port Balance Priority is not set
```

To rebalance the E_Port load on a switch:

```
switch:admin> dlsset --rebalance  
Route rebalance successful
```

To rebalance the E_Port load on all switches:

```
switch:admin> dlsset --rebalance -all  
Route rebalance all successful
```

SEE ALSO [aptPolicy](#), [dlsReset](#), [dlsShow](#), [iodReset](#), [iodSet](#), [iodShow](#), [uRouteShow](#), [topologyShow](#)

dlsShow

Displays the setting of Dynamic Load Sharing (DLS).

Synopsis `dlsshow`

Description Use this command to display information about Dynamic Load Sharing configuration settings on the switch. Depending on the configuration, the command output displays one of the following messages:

- **DLS is set with Lossless enabled** - DLS is enabled with the Lossless feature. Load sharing is recomputed with every change in the fabric, and existing routes can be moved to maintain optimal balance. In Lossless mode, no frames are lost during this operation.
- **DLS is set with Lossless disabled** - DLS is enabled without the Lossless feature. Load sharing is reconfigured with every change in the fabric, and existing routes can be moved to maintain optimal balance. No attempt is made to prevent frames from being lost while load sharing is recomputed.
- **DLS is not set with Lossless disabled** - DLS is disabled and the Lossless option is disabled. Existing routes are never moved to maintain optimal balance. If the Lossless option was enabled before you disabled DLS, it is now disabled as well. This means, frame loss is not prevented during a load sharing recomputation.
- **DLS is set with Two-hop Lossless enabled** - DLS is enabled with Lossless and two-hop lossless DLS.

Refer to **dlsSet** for a description of load sharing.

This command also displays the status of the E_Port Balance Priority feature on the switch. If the feature is enabled, the message "E-port Balance Priority is set" displays. If the feature is disabled, the message "E-port Balance Priority is not set" displays.

Notes The Lossless feature is not supported on GbE ports and FCoE ports. On unsupported platforms, this command exits with an appropriate message.

On platforms that do not support the E_Port Balance Priority feature, the command output indicates that the feature is not supported.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the current DLS and E_Port Balance Priority configuration on a switch with port-based routing:

```
switch:admin> dlsshow
DLS is set with Lossless enabled.
E_Port Balance Priority is set.
```

SEE ALSO [dlsSet](#), [dlsReset](#)

dnsConfig

Sets, displays, or removes domain name service (DNS) parameters.

Synopsis `dnsconfig`

`dnsconfig --add -domain name -serverip1 ipaddr serverip2 ipaddr`

`dnsconfig --delete`

`dnsconfig --show`

`dnsconfig --help`

Description

Use this command to display, set, or remove the domain name service parameters.

The domain name service parameters are the domain name and the name server IP address for primary and secondary name servers. The **dnsconfig** command displays IPv4 and IPv6 addresses.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

--add

Sets the DNS configuration.

-domain *name*

Specifies the DNS domain name.

-serverip1 *ipaddr*

Specifies the DNS primary server IP address. The IP address can be an IPv4 or IPv6 address.

-serverip2 *ipaddr*

Specifies the DNS secondary server IP address. The IP address can be an IPv4 or IPv6 address.

--delete

Removes the DNS configurations.

--show

Displays the current DNS configuration.

--help

Displays the command usage.

Examples

To set the DNS parameters for the system:

```
switch:admin> dnsconfig

Enter option
1 Display Domain Name Service (DNS) configuration
2 Set DNS configuration
```

```

3 Remove DNS configuration
4 Quit
Select an item: (1..4) [4] 2

Enter Domain Name: [] domain.com
Enter Name Server IP address in dot notation: [] \
123.123.123.123
Enter Name Server IP address in dot notation: [] \
123.123.123.124
DNS parameters saved successfully

Enter option
1 Display Domain Name Service (DNS) configuration
2 Set DNS configuration
3 Remove DNS configuration
4 Quit
Select an item: (1..4) [4] 4

```

To configure the DNS domain name, DNS server address:

```

switch:admin> dnsconfig --add -domain www.cp0.com \
serverip1 192.168.1.1 serverip2 192.168.201.1
DNS parameters saved successfully.

```

To delete the DNS configurations:

```

switch:admin> dnsconfig --delete
DNS parameters removed successfully.

```

To display the DNS configurations:

```

switch:admin> dnsconfig --show
Domain Name Server Configuration Information
-----
Domain Name           = www.cp0.com
Name Server IP Address = 192.168.1.1
Name Server IP Address = 192.168.201.1

```

SEE ALSO [configDownload](#), [configUpload](#), [firmwareDownload](#), [ipAddrSet](#), [ipAddrShow](#)

enclosureShow

Displays attributes of the switch enclosure.

Synopsis `enclosureshow attribute`

Description Use this command to display attributes of the switch enclosure, including the vendor-specific enclosure identifier and the identifier of the enclosure interface to which the switch is attached.

This command applies to products that are embedded in a blade server or storage chassis. Most options are platform-specific. Options that do not apply to a platform are identified with a "Not supported on this platform" message.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

attribute

Specifies the enclosure attribute. Valid attributes include the following:

id

The vendor-specific enclosure identifier.

modelName

The vendor-specific enclosure model name.

slotid

The identifier of the enclosure interface to which the switch is attached.

rackname

The name assigned by the enclosure manager to the rack.

rackid

The serial number assigned by the enclosure manager to the rack.

enclosurename

The name assigned by the enclosure manager to the enclosure.

enclosureid

The serial number assigned by the enclosure manager to the enclosure.

connname

The product name used by the enclosure manager for the switch model.

connaddr

The connector address used by the enclosure manager for this switch (indicates the physical position of the switch in the enclosure).

connid

The serial number of the switch used by the enclosure manager (not to be confused with the Factory Serial Number).

enclosureShow

| | |
|------------------|---|
| conntype | The connector type used by the enclosure manager for this model of switch. |
| connloc | The switch location within the enclosure. |
| connpres | Information about the presence of the switch that is used by the enclosure manager. |
| connfuse | Information about whether or not the switch has a fuse. |
| uuid | The Universal Unique ID for the switch if visible to the enclosure manager. |
| mamacaddr | The enclosure manager's Ethernet MAC Address. |
| snmpports | The SNMP agent and trap ports if visible to the enclosure manager. |

Examples To display the identifier of the enclosure interface to which the switch is attached:

```
switch:admin> enclosureShow slotid  
Bay 4
```

SEE ALSO [chassisShow](#)

errClear

Clears all error log messages for all switch instances on this Control Processor (CP).

Synopsis `errclear`

Description Use this command to clear all internal and external error log messages for all switch instances on the CP where the command is executed. For products with a single processor, all error log messages are cleared. For products that contain multiple processors, this command can be executed on either control processor. It clears the error log only on the CP where the command is executed. For example, to clear the error log on the standby CP, issue **errclear** on the standby CP.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To clear the error log messages:

```
switch:admin> errclear
```

SEE ALSO [errDump](#), [errShow](#)

errDelimiterSet

Sets the error log start and end delimiters for messages sent to the console and syslog.

Synopsis `errdelimiterset`
`[-s "start_delimiter_string"]`
`[-e "end_delimiter_string"]`

Description Use this command to set the error log start and end delimiters for log messages sent to the console and syslog. An empty string clears the start and the end delimiters (including the colon) so that they are not displayed.

When used without operands, this command displays the existing **errDelimiterSet** configuration. The delimiter configuration is stored persistently.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Start and end delimiters are displayed only at the console; they are not displayed in a Telnet session or in the RASLog.

Operands This command has the following operands:

`-s "start_delimiter_string"`
 Specifies an alphanumeric string for the start delimiter. The string can be up to 10 characters long and must be enclosed in double quotation marks. This operand is optional.

`-e "end_delimiter_string"`
 Specifies the alphanumeric string for the end delimiter. The string can be up to 10 characters long and must be enclosed in double quotation marks. This operand is optional.

Examples To display the start and end delimiters:

```
switch:admin> errdelimiterset
delimiter start string: <none>
delimiter end string: <none>
```

To change the start and end delimiters (with sample output):

```
switch:admin> errdelimiterset -s "Start" -e "End"

switch:admin> errdelimiterset
delimiter start string: Start
delimiter end string: End
```

Sample output:

```
Start2003/03/10-09:54:03, [NS-1002], 1035,, ERROR,
SWITCH43, Name Server received an invalid request
from device 10:02:32:A3:78:23:23:End
```

SEE ALSO [errDump](#), [errFilterSet](#), [errShow](#)

errDump

Displays the error log without pagination.

Synopsis

```
errdump [--all]
errdump [--count number]
errdump [--reverse]
errdump [--severity severity]
errdump [--slot slotnum]
errdump [--attribute attribute]
errdump [--message msgID]
errdump [--start YYYY/MM/DD-HH:MM:SS]
errdump [--stop YYYY/MM/DD-HH:MM:SS]
errdump [--help]
```

Description Use this command to dump external error log messages without any page breaks. When executed without operands, this command prints all error messages for the logical switch context in which the command is executed.

The output of this command is unique for each control processor (CP). On dual CP systems this command must be executed on each CP to obtain a complete record.

Specify both **--start** and **--stop** operands to display messages within a time period. Specify the **--start** operand alone to display messages from a specific time to the end of the log. Specify the **--stop** operand alone to display messages from the beginning of the log to the specified time.

The following information is displayed in each message:

| | |
|---------------------------------|---|
| Start delimiter | Delimiter string for the start of a message. |
| Timestamp | Timestamp for the message. |
| Message ID | Message identifier. |
| External sequence number | Sequence number for the message. |
| Security audit flag | Security audit message displayed as AUDIT. |
| Severity | Severity of the message. Valid values include INFO, WARNING, ERROR, and CRITICAL. |

| | |
|----------------------|--|
| Switch name | Switch name for the generator of this message, or "chassis". |
| Message | Message body. |
| End delimiter | Delimiter string for the end of a message. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|---|--|
| --all | Displays messages for the entire chassis for a user with chassis permissions. This operand is optional; if omitted, the messages for the current logical switch context are displayed. |
| --count <i>number</i> | Displays messages based on count value. |
| --reverse | Displays messages in reversed order. This operand is optional; if omitted, the messages display in chronological order. |
| --severity <i>severity</i> | Displays messages based on severity. |
| CRITICAL | Displays critical messages. |
| ERROR | Displays error messages. |
| WARNING | Displays warning messages. |
| INFO | Displays informational messages. |
| --slot <i>slotnum</i> | Displays messages based on slot. |
| --attribute <i>attribute</i> | Displays messages with the given attribute. Valid attributes values are ALL and FFDC. |
| --message <i>msgID</i> | Displays all raslog messages based on the message ID. |
| --start <i>YYYY/MM/DD-HH:MM:SS</i> | Displays messages from the given start time. The value of the time is related to the raslog storage time. |
| --stop <i>YYYY/MM/DD-HH:MM:SS</i> | Display messages that are logged up to the given end time. |

--help

Displays command usage.

Examples

To display the error log for the chassis:

```
switch:admin> errdump --all
Fabric OS: v8.0.1

2015/12/17-22:29:17, [SEC-1203], 9036, CHASSIS, INFO, \
Spir_67, Login information : Login successful via \
TELNET/SSH/RSH. IP Addr: 10.106.7.62

2015/12/17-22:29:17, [ZONE-1022], 9037, CHASSIS, INFO, \
Spir_67, The effective configuration has changed to meh.

2015/12/17-22:29:17, [FABR-1001], 9041, CHASSIS, WARNING, \
Spir_67, port 0, incompatible Long distance mode.

2015/12/17-22:29:17, [LOG-1000], 9043, CHASSIS, INFO, \
Spir_67, Previous message repeated 1 time(s)
(output truncated)
```

To display messages for a slot:

```
switch:admin> errdump --slot 4
Fabric OS: v8.0.1

2015/06/19-03:26:44, [HAM-1004], 31, SLOT 4 | CHASSIS, INFO, \
pluto_19, Processor rebooted - Reboot.

2015/06/19-03:26:44, [SULB-1003], 32, SLOT 4 | CHASSIS, INFO, \
pluto_19, Firmwarecommit has started.

2015/06/19-03:26:44, [IPAD-1001], 33, SLOT 4 | CHASSIS, INFO, \
pluto_19, CP/1 IPv6 manual fe80::224:38ff:fe1b:4400 DHCP Off.

2015/06/19-03:29:15, [IPAD-1000], 48, SLOT 4 | CHASSIS, INFO, \
pluto_19, CP/0 Ether/0 IPv6 autoconf
fd00:60:69bc:816:205:1eff:fe84:3f49/64 tentative DHCP Off.

2015/06/19-03:29:15, [IPAD-1000], 51, SLOT 4 | CHASSIS, INFO, \
pluto_19, CP/1 Ether/0 IPv6 autoconf
fd00:60:69bc:816:205:1eff:fe84:4aal/64 tentative DHCP Off.
```

To display messages based on severity:

```
switch:admin> errdump --severity WARNING
Fabric OS: v8.0.1

2015/06/18-20:24:52, [SULB-1001], 2, SLOT 5 | CHASSIS, WARNING, \
pluto_19, Firmwaredownload command has started.

2015/06/18-20:36:08, [FSSM-1003], 6, SLOT 5 | CHASSIS, WARNING, \
pluto_19, HA State out of sync.

2015/06/18-20:36:11, [FABR-1001], 8, SLOT 5 | FID 128, WARNING, \
sw0, port 268, ESC detected Fabric ID conflict with neighbor \
(FID 10).
```

```
2015/06/18-20:36:11, [FABR-1001], 9, SLOT 5 | FID 128, WARNING, \
sw0, port 281, ESC detected Fabric ID conflict with neighbor \
(FID 20).
```

```
2015/06/18-20:36:11, [FABR-1001], 10, SLOT 5 | FID 128, WARNING, \
sw0, port 265, ESC detected Fabric ID conflict with neighbor \
(FID 1).
```

To display messages based on attributes:

```
switch:admin> errdump --attribute FFDC
Fabric OS: v8.0.1
```

```
2015/06/19-03:58:42, [LOG-1001], 81, SLOT 5 | FFDC | CHASSIS, \
WARNING, pluto_19, A log message was dropped
```

To display messages based on count value:

```
switch:admin> errdump --count 3
Fabric OS: v8.0.1
```

```
2015/06/18-20:23:09, [LOG-1003], 1, SLOT 5 | CHASSIS, INFO, \
pluto_19, The log has been cleared.
```

```
2015/06/18-20:24:52, [SULB-1001], 2, SLOT 5 | CHASSIS, WARNING, \
pluto_19, Firmwaredownload command has started.
```

```
2015/06/18-20:24:52, [SULB-1036], 3, SLOT 5 | CHASSIS, INFO, \
pluto_19, The current Version: Fabric OS v8.0.1
```

To display raslog messages based on a message ID:

```
switch:admin> errdump --message SEC-1203
Fabric OS: v8.0.1
```

```
2015/11/21-11:06:21, [SEC-1203], 19, FID 128, INFO, switch, \
Login information: Login successful via TELNET/SSH/RSH. \
IP Addr: 10.70.12.104
```

```
2015/11/21-12:32:40, [SEC-1203], 20, FID 128, INFO, switch, \
Login information: Login successful via TELNET/SSH/RSH. \
IP Addr: 172.26.1.93
```

```
2015/11/22-10:27:41, [SEC-1203], 21, FID 128, INFO, switch, \
Login information: Login successful via TELNET/SSH/RSH. \
IP Addr: 10.70.4.109
```

To display messages from specified time to the end of the log:

```
switch:admin> errdump --start 2015/10/10-2042:06
Fabric OS: v8.0.1
```

```
2015/10/10-20:42:06, [XTUN-1997], 16, CHASSIS, WARNING, ch3456789012345678901234, \
FTRACE buffer 0 on slot 1 dp 2 has been triggered.
```

```
2015/10/10-20:42:06, [PLAT-1001], 17, CHASSIS, INFO, ch3456789012345678901234, \
CP0 resetting other CP (double reset may occur).
```

```
2015/10/10-20:42:06, [XTUN-1997], 18, CHASSIS, WARNING, ch3456789012345678901234, \
FTRACE buffer 0 on slot 1 dp 2 has been triggered.
```

```
2015/10/10-20:42:07, [PLAT-1001], 19, CHASSIS, INFO, ch3456789012345678901234, \
```

```

CP0 resetting other CP (double reset may occur).
2015/10/10-20:42:07, [XTUN-1997], 20, CHASSIS, WARNING, ch3456789012345678901234, \
FTRACE buffer 0 on slot 1 dp 2 has been triggered.
2015/10/10-20:42:07, [PLAT-1001], 21, CHASSIS, INFO, ch3456789012345678901234, \
CP0 resetting other CP (double reset may occur).

```

To display messages from beginning of the log to the specified time:

```

switch:admin> errdump --stop 2015/10/10-20:42:05
Fabric OS: v8.0.1
2015/10/10-20:41:58, [LOG-1003], 1, CHASSIS, INFO, ch3456789012345678901234, \
The log has been cleared.
2015/10/10-20:42:05, [XTUN-1997], 2, CHASSIS, WARNING, ch3456789012345678901234, \
FTRACE buffer 0 on slot 1 dp 2 has been triggered.
2015/10/10-20:42:05, [PLAT-1001], 3, CHASSIS, INFO, ch3456789012345678901234, \
CP0 resetting other CP (double reset may occur).
2015/10/10-20:42:05, [XTUN-1997], 4, CHASSIS, WARNING, ch3456789012345678901234, \
FTRACE buffer 0 on slot 1 dp 2 has been triggered.
2015/10/10-20:42:05, [PLAT-1001], 5, CHASSIS, INFO, ch3456789012345678901234, \
CP0 resetting other CP (double reset may occur).
2015/10/10-20:42:05, [XTUN-1997], 6, CHASSIS, WARNING, ch3456789012345678901234, \
FTRACE buffer 0 on slot 1 dp 2 has been triggered.
2015/10/10-20:42:05, [PLAT-1001], 7, CHASSIS, INFO, ch3456789012345678901234, \
CP0 resetting other CP (double reset may occur).

```

To display messages within a time period:

```

switch:admin> errdump --start 2015/10/10-20:42:06 --stop 2015/10/10-20:42:07
Fabric OS: v8.0.1
2015/10/10-20:42:06, [XTUN-1997], 8, CHASSIS, WARNING, \
ch3456789012345678901234, FTRACE buffer 0 on slot 1 \
dp 2 has been triggered.
2015/10/10-20:42:06, [PLAT-1001], 9, CHASSIS, INFO, \
ch3456789012345678901234, CP0 resetting other CP \
(double reset may occur).
2015/10/10-20:42:06, [XTUN-1997], 10, CHASSIS, WARNING, \
ch3456789012345678901234, FTRACE buffer 0 on slot 1 \
dp 2 has been triggered.
2015/10/10-20:42:07, [PLAT-1001], 19, CHASSIS, INFO, \
ch3456789012345678901234, CP0 resetting other CP \
(double reset may occur).
2015/10/10-20:42:07, [XTUN-1997], 20, CHASSIS, WARNING, \
ch3456789012345678901234, FTRACE buffer 0 on slot 1 \
dp 2 has been triggered.
2015/10/10-20:42:07, [PLAT-1001], 21, CHASSIS, INFO, \
ch3456789012345678901234, CP0 resetting other CP \
(double reset may occur).

```

SEE ALSO [errDelimiterSet](#), [errFilterSet](#), [errShow](#)

errFilterSet

Sets a filter for an error log destination.

Synopsis `errfilterset [-d destination][-v severity]`

Description Use this command to set a filter for an error log destination. A filter is set based on the severity level of the messages.

When used without operands, this command displays the filters that are currently in configured.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-d *destination* Specifies the destination for the filter. The string **console** is the only valid value at this time.

-v *severity* Specifies the minimum severity level of the message to pass through the filter. Valid values are INFO, WARNING, ERROR, or CRITICAL. Input values are not case-sensitive.

Examples To display the current filter settings:

```
switch:admin> errfilterset
console: filter severity = WARNING
```

To set the filter severity level for the console:

```
switch:admin> errfilterset -d console -v warning
```

SEE ALSO [errDump](#), [errShow](#)

errModuleShow

Displays all the defined error log modules.

Synopsis `errmoduleshow`

Description Use this command to display a list of all defined error log modules.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display a list of all defined error log modules:

```
switch:user> errmoduleshow
Module IDs:

    1 KT          2 UT          3 TRCE        4 KTRC
    5 LOG         6 CDR         7 BLPU        8 PISP

    9 PIXE       10 EGR        11 BL         12 PIC
   13 PS        14 RTE        15 AS         16 AUTH
   17 BLDE      18 BLM        19 BPRT       20 CER
   21 CFLD      22 CFMN       23 CHPS       24 CONF
(output truncated)
```

SEE ALSO [errDump](#), [errShow](#)

errShow

Displays the error log messages with pagination.

Synopsis `errshow [-a | -r | -t | -o]`

Description Use this command to display external error log messages one at a time. When executed without operands, this command prints the error messages for the logical switch context in which the command is executed. When used with the **-a** option, the command prints the error messages for the entire chassis. The messages are displayed with page breaks. The **-r** operand displays the messages in reversed order.

The output of this command is unique for each Control Processor (CP). On dual CPs this command must be executed on each CP to obtain a complete record.

Specify both **-t** and **-o** operands to display messages within a time period. Specify the **-t** operand alone to display messages from a specific time to the end of the log. Specify the **-o** operand alone to display messages from the beginning of the log to the specified time.

The following information displays in each message:

| | |
|---------------------------------|---|
| Start delimiter | Delimiter string for the start of a message. |
| Timestamp | Timestamp for the message. |
| Message ID | Message identifier. |
| External sequence number | Sequence number for the message |
| Security audit flag | Security audit message displayed as AUDIT. |
| Severity | Severity of the message. Valid values include INFO, WARNING, ERROR, and CRITICAL. |
| Switch name | Switch name for the generator of this message, or "chassis". |
| Message | Message body. |
| End delimiter | Delimiter string for the end of a message. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

- a** Displays messages for the entire chassis for a user with admin and chassis permissions. This operand is optional; if omitted, the messages for the current logical switch context are displayed.
- r** Displays messages in reversed order. This operand is optional; if omitted, the messages display in the chronological order.
- t** Displays messages from the given start time. The value of the time is related to the raslog storage time.
- o** Display messages that are logged up to the given end time.

Examples To display the error log for the chassis:

```
switch:admin> errshow -a
Fabric OS: v8.0.1

2015/06/11-10:16:07, [SEC-1203], 261, FID 128, INFO, \
switch, Login information: Login successful via \
TELNET/SSH/RSH. IP Addr: 192.0.2.2

[Type <CR> to continue, Q<CR> to stop:

2015/06/11-10:44:13, [TS-1002], 262, FID 128, INFO, \
switch, External Clock Server used instead of \
LOCL: loc1: 0x45585400 remote: 0x4c4f434c.

[Type <CR> to continue, Q<CR> to stop:

2015/06/11-11:00:49, [SNMP-1005], 287, FID 128, INFO, \
switch, SNMP configuration attribute, SNMPv3 Trap \
Recipient IP Address 6, has changed from [192.0.2.10] \
to [0.0.0.0].

[Type <CR> to continue, Q<CR> to stop:

2015/06/11-11:00:49, [SNMP-1005], 288, FID 128, INFO, \
switch, SNMP configuration attribute, SNMPv3 Trap \
Recipient Severity Level 6, has changed from 4 to 0.

[Type <CR> to continue, Q<CR> to stop:
```

To display messages from specified time to the end of the log:

```
switch:admin> errshow -t 2016/10/31-19:17:23
Fabric OS: v8.0.0
2015/10/31-19:17:25, [IPAD-1000], 15, CHASSIS, INFO, Brocade5300, \
  SW/0 Ether/0 IPv6 autoconf 2012::90/64 DHCP Off.

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:17:25, [HAM-1004], 16, CHASSIS, INFO, Brocade5300, \
  Processor rebooted - Hareboot.

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:17:29, [FV-1001], 17, CHASSIS, INFO, switch, \
  Flow Vision daemon initialized.

Type <CR> to continue, Q<CR> to stop:
```

To display messages from beginning of the log to the specified time:

```
switch:admin> errshow -o 2015/10/31-19:01:05
Fabric OS: v8.0.1
2015/10/31-18:59:52, [LOG-1003], 1, CHASSIS, INFO, Brocade5300, \
  The log has been cleared.

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:00:09, [SEC-1203], 2, FID 128, INFO, switch, \
  Login information: Login successful via TELNET/SSH/RSH. IP Addr: 10.70.12.103

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:00:10, [SEC-1203], 3, FID 128, INFO, switch, \
  Login information: Login successful via TELNET/SSH/RSH. IP Addr: 10.70.12.103

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:00:16, [RAS-1007], 4, CHASSIS, INFO, Brocade5300, \
  System is about to reload.

Type <CR> to continue, Q<CR> to stop:
```

To display messages within a time period:

```
switch:admin> errshow -t 2015/10/31-19:01:06 -o 2015/10/31-19:01:11
Fabric OS: v8.0.1
2015/10/31-19:01:06, [IPAD-1000], 5, CHASSIS, INFO, Brocade5300, \
  SW/0 Ether/0 IPv6 autoconf 2012::90/64 DHCP Off.

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:01:06, [IPAD-1001], 6, CHASSIS, INFO, Brocade5300, \
  CP/0 IPv6 autoconf fe80::224:38ff:fea5:6200 DHCP not Set.

Type <CR> to continue, Q<CR> to stop:

2015/10/31-19:01:06, [IPAD-1000], 7, CHASSIS, INFO, Brocade5300, \
  SW/0 Ether/0 IPv6 autoconf 2620:100:4:f801:205:33ff:fe00:9b28/64 tentative DHCP
```

Off.

Type <CR> to continue, Q<CR> to stop:

```
2015/10/31-19:01:06, [HAM-1004], 8, CHASSIS, INFO, Brocade5300, \  
Processor rebooted - Fastboot.
```

Type <CR> to continue, Q<CR> to stop:

```
2015/10/31-19:01:06, [IPAD-1000], 9, CHASSIS, INFO, Brocade5300, \  
SW/0 Ether/0 IPv6 autoconf 2620:100:4:f801:205:33ff:fe00:9b28/64 DHCP Off.
```

Type <CR> to continue, Q<CR> to stop:

```
2015/10/31-19:01:11, [FV-1001], 10, CHASSIS, INFO, switch, \  
Flow Vision daemon initialized.
```

Type <CR> to continue, Q<CR> to stop:

SEE ALSO [errDelimiterSet](#), [errDump](#), [errFilterSet](#)

ethIf

Displays and sets the link operating mode for a network interface.

Synopsis `ethif --set interface -auto-negotiate | -an [on | off] -speed [speed] -duplex [full | half]`

`ethif --reseterror interface`

`ethif --show interface`

`ethif --help`

Description Use this command to set and view the link operating mode for a network interface.

Changing the link mode is not supported for all network interfaces or for all Ethernet network interfaces. On the CP of a Brocade DCX or DCX-4S, this command supports **eth0** and **eth3** as interface parameters. On all other platforms, only **eth0** is supported.

The CP on a Brocade DCX or DCX-4S has two external physical Ethernet management ports, **eth0** and **eth3**. Both interfaces are bonded together to form a single logical interface, **bond0**.

The management port IP addresses are assigned to the logical interface, **bond0**. Link layer Ethernet operations are applied to the physical interfaces, **eth0** and **eth3**.

Ethernet bonding provides link layer redundancy using the active-standby failover model. The two Ethernet ports must be part of the same subnet. By default, all traffic is transmitted over the active Ethernet port, **eth0**. The second Ethernet port, **eth3**, acts as a standby interface and no traffic is transmitted over it. When the active Ethernet port is disconnected, the alternate Ethernet port becomes active. When the system reboots, the Ethernet port **eth0** is always made active if it is connected.

When executed with the **bond0** operand, **ethif --show** displays the active Ethernet port.

When selecting autonegotiation, you can choose the specific link operating modes that are advertised to the link partner. At least one common link operating mode must be advertised by both sides of the link.

When forcing the link operating mode, both sides of the link must be forced to the same mode. The link does not work reliably if one side is set to autonegotiate and the other side is set to forced mode.

Exercise care when using this command. Forcing the link to an operating mode not supported by the network equipment to which it is attached may result in an inability to communicate with the system through its Ethernet interface. It is recommended that this command be used only from the serial console port. When used through an interface other than the serial console port, the command displays a warning message and prompts for verification before continuing. This warning is not displayed and you are not prompted when the command is used through the serial console port.

For dual-CP systems, this command affects only the CP to which you are currently logged in to set the link operating mode on the active CP; you must issue this command on the active CP; to set the link operating mode on the standby CP, you must issue this command on the standby CP. During failover, the link operating mode is retained separately for each CP, because the physical links may be set to operate in different modes.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|-------------------------------|--|
| --set | Sets a link operating mode for a network interface. |
| <i>interface</i> | Specifies the name of the interface, optionally enclosed in double quotation marks. Valid interfaces include the following: |
| eth0 eth3 | Displays the link operating mode of the specified interface. The value eth3 is valid only on the CP of a Brocade DCX or DCX-4S. |
| bond0 | Displays the active Ethernet port. This operand is valid only on the CP of a DCX or a DCX-4S. |
| -auto-negotiate -an | Configures auto-negotiation. |
| on off | Sets auto-negotiation on or off. |
| -speed speed | Configures the speed. You can configure 1000, 100, or 10 as the speed. Note that 1000 Mbps speed can be configured only when auto-negotiation is on. |
| -duplex | Configures whether the capability of operation is full duplex or half duplex. |
| --reseterror interface | Resets error counters on the network interface. The valid values for <i>interface</i> are eth0 and eth3 . |
| --show | Displays the link operating mode for a network interface. |
| <i>interface</i> | Specifies the name of the interface, optionally enclosed in double quotation marks. Valid interfaces include the following: |
| eth0 eth3 | Displays the link operating mode of the specified interface. The value eth3 is valid only on the CP of a Brocade DCX or DCX-4S. |
| bond0 | Displays the active Ethernet port. This operand is valid only on the CP of a DCX or a DCX-4S. |
| --help | Displays the command usage. |

Examples To reset the link operating attributes on eth0:

```
switch:admin> ethif --reseterror eth0
Statistics cleared for eth0
```

To view the link operating attributes on bond0:

```
switch:admin> ethif --show bond0
bond0 interface:

bond0 includes physical interfaces: eth0 eth3
Currently Active Interface: eth0
Currently Slave Interface: eth3
Link mode: negotiated 1000baseT-FD, link ok
MAC Address: 00:05:1E:86:7B:79

bond0      Link encap:Ethernet  HWaddr 00:05:1E:86:7B:79
           inet addr:10.38.18.135  Bcast:10.38.31.255  Mask:255.255.240.0
           UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
           RX packets:2373 errors:0 dropped:0 overruns:0 frame:0
           TX packets:56 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:0
```

To view the link operating attributes on a network interface:

```
switch:admin> ethif --show eth0
eth0 interface:

Link mode: negotiated 1000baseT-FD, link ok
MAC Address: 00:05:1E:86:7B:79

eth0      Link encap:Ethernet  HWaddr 00:05:1E:86:7B:79
           UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
           RX packets:2397 errors:0 dropped:0 overruns:0 frame:0
           TX packets:55 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:1000
```

To set the link operating attributes on a network interface:

```
switch:admin> ethif --set eth0 -an on -speed 100 -duplex full
an:on
speed:100
cap:full
MII_CMD:-A

ADVERTISE:Advertise
DEFMODE:yes
auto:1
MII_MODE:100baseTx-FD,
Committing configuration...done.
```


To check the link operating attributes:

```
switch:admin> ethif --show eth0
Link mode: negotiated 100baseTx-FD, link ok
MAC Address: 00:05:1E:86:7B:79

eth0      Link encap:Ethernet  HWaddr 00:05:1E:86:7B:79

          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:2397 errors:0 dropped:0 overruns:0 frame:0
          TX packets:56 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
```

SEE ALSO [ipAddrSet](#), [ipAddrShow](#)

exit

exit

Description See [logout](#).

SEE ALSO None

extnCfg

Sets the extension mode configuration.

Synopsis `extncfg --ve | --ve-mode [--slot slot] 10VE | 20VE`

`extncfg --app-mode fcip | hybrid`

`extncfg --config -default | -clear [-slot slot | -all]`

`extncfg --show`

`extncfg --help`

Description Use this command to set the extension mode configuration. This command switches the Field Programmable Gate Array (FPGA) images or the VE_Port mappings.

Notes The execution of this command is disruptive and requires the switch reboot.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

`--slot slot/`

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

`--ve | --ve-mode 10VE | 20VE`

Switches the VE_Port mode between 10 VE-ports and 20 VE-ports. The default VE_Port mode is 10VE.

`--app-mode fcip | hybrid`

Sets the system extension mode to FCIP or hybrid (FCIP with IPEXT). The default APP_Port mode is FCIP. Configuring the switch for hybrid mode is disruptive. The switch reboots and loads the hybrid image.

`--config -default`

Removes active extension configurations from the specified slot or all slots.

`--config -clear`

Removes the inactive (stale) extension configurations from the specified slot or all slots. After executing this command, you must run the **slotPowerOff** and **slotPowerOn** commands for the extension blade to clear the configurations completely.

`-all`

For bladed systems only, indicates all slots in the chassis.

`-show`

Displays the current extension mode configuration.

`--help`

Displays the command usage.

Examples To display the current extension mode configuration:

```
switch:admin> extncfg --show
App Mode is FCIP
VE-Mode: configured for 10VE mode.
```

To switch the extension mode configuration:

```
switch:admin> extncfg --ve-mode 20VE
This operation will require a reboot of the switch and will disrupt
any running traffic.
Are you sure you want to proceed? (y/n): y

Operation Succeeded. Rebooting Switch.
```

To configure the system for hybrid mode:

```
switch:admin> extncfg --app-mode hybrid
This action will configure the system for Hybrid (FCIP/IPEXT) mode.

WARNING: This is a disruptive operation that requires a reboot \
to take effect. Would you like to continue (Y,y,N,n): [ n ] y

Operation succeeded. Rebooting the system...

switch:admin> extncfg --show
APP Mode is Hybrid (FCIP with IPEXT)
VE-Mode: configured for 20VE mode.
```

SEE ALSO [Fabric OS Commands](#)

fabRetryShow

Displays the retry count of the fabric commands.

Synopsis `fabretryshow`

Description Use this command to display the retry count of the fabric commands. For each port, the command output displays counts for the following Switch Internal Link Service (SW_ILS) requests:

| | |
|----------------|---|
| ELP | Exchange Link Parameters |
| EFP | Exchange Fabric Parameters |
| HA_EFP | Exchange Fabric Parameters used during warm recovery |
| DIA | Domain Identifier Assigned |
| RDI | Request Domain Identifier |
| BF | Build Fabric |
| FWD | Fabric Controller Forward |
| EMT | Fabric Controller Mark Timestamp |
| ETP | Fabric Controller Trunk Parameters |
| RAID | Return Address Identifier |
| GAID | Get Address Identifier |
| ELP_TMR | Used internally for fabric application (not a SW_ILS) |
| GRE | Get Route Entry |
| ECP | Exchange Credit Parameters |
| ESC | Exchange Switch Capabilities |
| EFMD | Exchange Fabric Membership Data |

fabRetryShow

ESA Exchange Security Attributes

DIAG_CMD Diagnostic Command

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the retry count of Fabric OS Commands:

```
switch:user> fabretryshow
                SW_ILS
E_Port  ELP  EFP  HA_EFP  DIA  RDI  BF   FWD  EMT  ETP  RAID  \
34      0   0   0       0   0   0   0   0   0   0   \
35      0   0   0       0   0   0   0   0   0   0   \

GAID  ELP_TMR  GRE  ECP  ESC  EFMD  ESA  DIAG_CMD
0     0       0   0   0   0   0   0
0     0       0   0   0   0   0   0
```

SEE ALSO None

fabRetryStats

Displays or manages the retry count of fabric commands.

Synopsis `fabretrystats --show`

`fabretrystats --clear`

`fabretrystats --help`

Description Use this command to view and clear the retry count of the fabric commands. The `--clear` option clears the counters for all the Switch Internal Link Service (SW_ILS) requests. For each E/D_Port, the `--show` option displays the counters for the following Switch Internal Link Service (SW_ILS) requests:

| | |
|----------------|---|
| ELP | Exchange Link Parameters |
| EFP | Exchange Fabric Parameters |
| HA_EFP | Exchange Fabric Parameters used during warm recovery |
| DIA | Domain Identifier Assigned |
| RDI | Request Domain Identifier |
| BF | Build Fabric |
| FWD | Fabric Controller Forward |
| EMT | Fabric Controller Mark Timestamp |
| ETP | Fabric Controller Trunk Parameters |
| RAID | Return Address Identifier |
| GAID | Get Address Identifier |
| ELP_TMR | Used internally for fabric application (not a SW_ISL) |
| GRE | Get Route Entry |
| ECP | Exchange Credit Parameters |

| | |
|-----------------|---------------------------------|
| ESC | Exchange Switch Capabilities |
| EFMD | Exchange Fabric Membership Data |
| ESA | Exchange Security Attributes |
| DIAG_CMD | Diagnostic Command |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

- show** Displays the retry count of fabric commands for all E_Ports and D_Ports.
- clear** Clears the retry count of fabric commands.
- help** Displays the command usage.

Examples To display the retry count of fabric commands:

```
switch:admin> fabretrystats --show
                        SW_ILS
-----\
E/D_Port  ELP  EFP  HA_EFP  DIA  RDI  BF  FWD  EMT \
-----\
3          1    0    0       0    0    0    0    0 \
12         0    2    0       0    0    0    0    0 \
      ETP  RAID  GAID  ELP_TMR  GRE  ECP  ESC  EFMD  ESA  DIAG_CMD
      0    0    0    0       0    0    0    0    0    0
      0    0    0    0       0    0    0    0    0    0
```

To clear the retry count of fabric commands:

```
switch:admin> fabretrystats --clear
switch:admin> fabretrystats --show
                        SW_ILS
-----\
E/D_Port  ELP  EFP  HA_EFP  DIA  RDI  BF  FWD  EMT \
-----\
3          0    0    0       0    0    0    0    0 \
12         0    0    0       0    0    0    0    0 \
      ETP  RAID  GAID  ELP_TMR  GRE  ECP  ESC  EFMD  ESA  DIAG_CMD
      0    0    0    0       0    0    0    0    0    0
      0    0    0    0       0    0    0    0    0    0
```

SEE ALSO [fabRetryShow](#), [fabStatsShow](#)

fabricLog

Displays (all users) or manipulates (admin) the fabric log.

Synopsis `fabriclog -s | --show [dport]`

`fabriclog -c | --clear [dport]`

`fabriclog -d | --disable`

`fabriclog -e | --enable`

`fabriclog -h | --help`

Description

Use this command to display, clear, disable, or enable the fabric log. When used with the `--show` option, this command displays the following information:

Time Stamp

Time of the event in the following format *HH:MM:SS:MS*.

Input and *Action

Fabric log message. An asterisk (*) in the message indicates an action. The link reset information is indicated by LR_IN and LR_OUT. LR_IN indicates a link reset on the remote switch, whereas LR_OUT indicates a link reset on the local switch.

S

Current switch state. Valid switch states include the following:

FO

Build Fabric (BF) received.

F1

Reconfigure Fabric (RCF) is not supported.

F2

Exchange Fabric Parameters (EFP) is waiting for last Accept Frame (ACC) from flood.

F3

Flood EFPs.

D0

The switch is the principal switch.

D1

The principal switch is sending Domain ID Assigned (DIA).

D2

The principal switch is waiting for Request Domain ID (RDI).

D3

The principle switch is processing the RDI.

A0

The switch is not the principal switch.

| | |
|-----|--|
| A1 | The non-principal switch is sending a DIA. |
| A2 | The non-principal switch is waiting for an RDI. |
| A3 | The non-Principal switch is processing an RDI. |
| S0 | The switch is in offline state. |
| P | Port state. Port states include the following: |
| PO | The port is offline. |
| P1 | The port is online. |
| P2 | Exchange Link Parameters (ELP) Accept Frame (ACC) received. |
| P3 | Link reset occurred on master or E_Port. |
| I0 | Trunk Initiator: Exchange Mark Timestamp (EMT) sent. |
| I1 | Trunk Initiator: Exchange Trunking Parameters (ETP) Accept Frame (ACC) received. |
| I2 | Trunk Initiator: ETP sent. |
| I3 | Trunk Initiator: Link reset occurred. |
| I4 | Trunk Initiator: Link reset done on slave. |
| T0 | Trunk Target: EMT received. |
| T1 | Trunk Target: ETP received. |
| T2 | Trunk Target: Link reset. |
| T3 | Trunk Target: Link reset done on slave. |
| LD | Dynamic long distance ECP sent or received. |
| ESC | Exchange Switch Capabilities (ESC) state between P2 and P3. |

| | |
|-------------|---|
| Sn | Next switch state. Refer to switch states for valid states. |
| Pn | Next port state. Refer to port states for valid states. |
| Port | Port number, if applicable, or NA |
| Xid | OXID (Exchange ID) Port number, if applicable, or NA |

Refer to the FC-SW-5 specification for more information on the port and switch states logged by this command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|-----------------------------|--|
| -s --show [dport] | Displays the fabric log. Optionally displays the D_Port logs only. |
| -c --clear [dport] | Clears the fabric log. Optionally clears the D_Port logs only. |
| -d --disable | Disables the fabric log. By default, the fabric log is enabled. |
| -e --enable | Enables the fabric log. |
| -h --help | Displays the command usage. |

Examples To display the fabric log:

```
switch:admin> fabriclog -s
Time Stamp      Input and *Action                                     S, P   Sn,Pn  Port  Xid
=====
Switch 0; Thu Feb 23 06:55:59 2012 GMT (GMT+0:00)
06:55:59.661357 *Fss Init                             NA,NA  NA,NA  NA    NA
06:55:59.661606 *Initiate State                       NA,NA  F2,NA  NA    NA
06:55:59.963652 Expd1 0x00000000 0000ffff ffffffff ffffffff F2,NA  F2,NA  0     NA
06:56:03.242214 Rcv FSS_RECOV_COLD                    F2,NA  F2,NA  NA    NA
06:56:03.242264 D-port Offline Skip Cnt 1(inst = 1)   F2,NA  F2,NA  NA    NA
06:58:44.880675 SCN LR_PORT (0);g=0x22; LR_IN        A2,P0  A2,P0  41    NA
07:00:12.107354 D-port Offline Skip Cnt 1(inst = 4)   F2,NA  F2,NA  NA    NA
07:00:12.640790 SCN LR_PORT (0);g=0x0; LR_OUT        F2,P0  F2,P0  11    NA
07:00:12.772930 SCN Port Online;g=0x0;                F2,P0  F2,P1  11    NA
(output truncated)
```

fabricLog

To display the cleared fabric log:

```
switch:admin> fabriclog -s
Time Stamp  Input and *Action  S, P  Sn,Pn  Port  Xid
=====
Number of entries: 0
Max number of entries: 64
```

SEE ALSO None

fabricName

Configures the fabric name parameter.

Synopsis `fabricname --set fabric_name`

`fabricname --clear`

`fabricname --show`

`fabricname --help`

Description Use this command to configure a name for the fabric.

With Virtual Fabrics it is not uncommon to have multiple fabrics in a single chassis. These logical fabrics are identified by their Fabric ID. With the **fabricName** command, you can address a fabric by name. This command provides an option for addressing fabrics by name. However, the fabric name does not replace the FID; the functionality of the FID is unaffected by the fabric name.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|---------------------------------------|--|
| <code>--set <i>fabric_name</i></code> | Sets the fabric name. The fabric name is a unique character string and can be up to 128 characters long. You can use any character including spaces and meta-characters. A fabric name that includes spaces must be included in double quotation marks. Tabs are not permitted. By default, the fabric name is not configured. |
| <code>--clear</code> | Clears the fabric name. |
| <code>--show</code> | Displays the fabric name. If a fabric name is not configured, an appropriate message is displayed. |
| <code>--help</code> | Displays the command usage. |

Examples To set and display the fabric name:

```
switch:user> fabricname --show
Fabric Name is not configured

switch:user> fabricname --set myfabric@1
Fabric Name set to "myfabric@1"

switch:user> fabricname --show
Fabric Name: "myfabric@1"
```

fabricName

To set a fabric name that includes spaces:

```
switch:user> fabricname --set "my new fabric"  
Fabric Name set to "my new fabric"
```

```
switch:user> fabricname --show  
Fabric Name: "my new fabric"
```

To clear the fabric name:

```
switch:user> fabricname --clear  
Fabric Name cleared!
```

```
switch:user> fabricname --show  
Fabric Name is not configured
```

SEE ALSO [switchShow](#), [fabricShow](#)

fabricPrincipal

Sets the principal switch selection mode.

Synopsis `fabricprincipal --help | -h`

`fabricprincipal [--show | -q]`

`fabricprincipal --enable`
`[-priority | -p priority]`
`[-force | -f]`

`fabricprincipal --disable`

`fabricprincipal [-f] mode`

Description Use this command to set principal switch selection mode for a switch and to set priorities for principal switch selection.

The implementation of the **fabricPrincipal** command is based solely on mechanisms specified in the Fibre Channel standards. These mechanisms provide a preference for a switch requesting to be the principal switch in a fabric, but they do not provide an absolute guarantee that a switch requesting to be the principal switch is granted this status.

When dealing with large fabrics, the selection of the principal switch is less deterministic. In these cases, to help ensure that the desired switch is selected as the principal switch, a small cluster of switches should be interconnected first, followed by additional switches to enlarge the fabric.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--help | -h Displays the command usage.

--show Displays the current mode setting and principal switch selection priority. This operand is optional; if not specified, **fabricPrincipal** displays the same data as with the **--show** option.

-q Displays principal mode only (enabled or disabled). This is a legacy command option that does not display the priority settings.

--enable Enables principal switch selection. The following operands are optional. If you do not provide a priority value, the system assigns the default of 0x01 or generates a value based on the switch state.

-priority | -p *priority* Sets the principal selection priority of the switch. The specified priority value is used in the principal switch selection protocol when the fabric rebuilds. Not all of these values can be assigned.

fabricPrincipal

| | |
|--------------------|--|
| 0x00 | Reserved. This value cannot be assigned. |
| 0x01 | Highest priority. This is a user-defined value |
| 0x02 | Switch was principal prior to sending or receiving a build fabric (BF) request. This value is generated by the switch to initiate a fabric reconfiguration. This value should not be assigned. |
| 0x3 - 0xFE | Priority value range. Choose a value in this range to indicate priority. Higher numbers mean lower priority. |
| 0xFF | Switch is not capable of acting as a principal switch. This is a user-defined value. Use --enable with a new priority to revert to this condition. |
| -force -f | Forces a fabric rebuild regardless of whether the switch is principal or subordinate. This option is not valid with the --disable command. |
| --disable | Disables principal switch selection. This command resets the priority to the default value 0xFE. |
| [-f] mode | Sets the principal switch selection mode. Specify 1 to enable principal switch selection mode. Specify 0 to disable principal switch selection mode. Optionally, use the -f operand to force a fabric rebuild. Mode changes take effect when the fabric rebuilds. This operand is optional. |

Examples

To enable a high fabric principal priority setting:

```
switch:admin> fabricprincipal --enable -p 0xff
Principal Selection Mode enabled (Activate in \
next fabric rebuild)
```

To disable the principal mode selection:

```
switch:admin> fabricprincipal --disable
Principal Selection Mode disabled
```

To display the current mode setting:

```
switch:admin> fabricprincipal -q
Principal Selection Mode: Enable
```

To disable the mode setting:

```
switch:admin> fabricprincipal 0
Principal Selection Mode disabled
```


To enable the mode setting:

```
switch:admin> fabricprincipal 1  
Principal Selection Mode enabled
```

To enable the mode setting and force fabric rebuild:

```
switch:admin> fabricprincipal -f 1  
Principal Selection Mode enabled \  
(Forcing fabric rebuild)
```

To display the principal switch selection priority:

```
switch:admin> fabricprincipal --show  
Principal Selection Mode: Enable  
Principal Switch Selection Priority: 0x10
```

SEE ALSO [fabricShow](#)

fabricShow

Displays fabric membership information.

Synopsis `fabricshow [-membership | -chassis | -paths | -version]`

`fabricshow -help`

Description Use this command to display information about switches in the fabric.

If the switch is initializing or is disabled, the message "no fabric" is displayed.

Running this command on an FCR or edge switch does not provide any router information; running this command on an edge switch with the **-membership** option does provide router information.

If the fabric is reconfiguring, some or all switches may not be displayed; otherwise, the following fields are displayed depending on the command option used:

| | |
|--------------------------|--|
| Switch ID | The switch Domain_ID and embedded port D_ID. |
| Worldwide Name | The switch WWN. |
| Enet IP Addr | The switch Ethernet IP address for IPv4- and IPv6-configured switches. For IPv6 switches, only the static IP address displays. |
| FC IP Addr | The switch IP FC (IP over Fibre Channel) address. |
| Name | The switch symbolic name. An arrow (>) indicates the principal switch. |
| FC Router IP Addr | The IP address of the FC Router. This field is empty if the switch is not an FC Router or it the FC Router does not support it. |
| FC Router Name | The FC Router symbolic name. This field is empty if the switch is not an FC Router or it the FC Router does not support it. |
| Chassis WWN | The world wide name of the chassis. For switches running firmware versions that do not distribute the chassis WWN, this field displays "NA". |
| Chassis Name | The name of the chassis. For switches running firmware versions that do not distribute the chassis name, this field displays "NA". |
| Fabric Name | If the fabric has an assigned name, the fabric name is displayed at the end of the command output. |

Path Count The number of currently available paths to the remote domain.

Version The firmware version of the domain.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-membership Displays fabric membership information with additional details of the FC Router, if present in the fabric.

-chassis Displays information about the chassis including chassis WWN and chassis name.

-paths Displays the number of paths available to each remote domain.

-version Displays firmware version details for each domain.

-help Displays the command usage.

Examples The following example illustrates a fabric of four switches. The switch named "sw180" is the Principal switch. Three of the switches are configured to run IP over Fibre Channel. The fabric name is configured.

```
switch:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   FC IP Addr   Name
-----
64:fffc40  10:00:00:60:69:00:06:56  192.168.64.59  192.168.65.59  "sw5"
65:fffc41  10:00:00:60:69:00:02:0b  192.168.64.180 192.168.65.180> "sw180"

66:fffc42  10:00:00:60:69:00:05:91  192.168.64.60  192.168.65.60  "sw60"
67:fffc43  10:00:00:60:69:10:60:1f  192.168.64.187  0.0.0.0        "sw187"

The Fabric has 4 switches
Fabric Name: mainFabricA
```

To show a mixed fabric with IPv4- and IPv6-configured switches (the fabric name is not configured):

```
switch:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr   FC IP Addr   Name
-----
1:fffc41  10:00:00:60:69:00:02:0b  192.168.64.180 192.168.65.180> "sw180"
                                     1080::8:800:200C:1234/64
2:fffc42  10:00:00:60:69:00:05:91  192.168.64.60  192.168.65.60  "sw60"

The Fabric has 2 switches.
```

To show additional details of the FC Router, if present:

```
switch:admin> fabricshow -membership
Switch ID      Name                ENET IP Addr      FC Router
                IP Addr          IP Addr          Name
-----
  1: fffc01  fcr_sprint_01  10.33.59.224
160: fffca0  fcr_fd_160     0.0.0.0           10.33.59.25  fcr_meteor2
190: fffcbe  fcr_mojo_6     10.33.59.32

The Fabric has 3 switches
```

To show additional details about the chassis:

```
switch:admin> fabricshow -chassis
Switch ID      Name                ENET IP Addr      Chassis WWN      Chassis Name
-----
  4:fffc04  sw5100_126_128  10.38.17.126  10:00:00:05:1e:0e:eb:58  Brcd5100
  5:fffc05  sw1500_127_128  10.38.17.127  10:00:00:05:1e:0e:eb:98  Brcd5100

The Fabric has 2 switches
```

To show additional details about the paths:

```
switch:admin> fabricshow -paths
Switch ID      Name                ENET IP Addr      Worldwide Name      Path Count
-----
  28: fffc02  DCX_28            192.0.2.28        10:00:00:05:1e:e5:ab:00  N/A
  38: fffc03  DCX_38            192.0.2.38        10:00:00:05:1e:47:cb:00   9

The Fabric has 2 switches
```

To show firmware version details:

```
switch:admin> fabricshow -version
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Version
-----
  1: fffc01  10:00:00:05:1e:82:3c:2a  10.25.224.140    11.22.33.44     v7.3.0
```

SEE ALSO [fabricName](#), [switchShow](#)

fabStatsShow

Displays fabric statistics.

Synopsis `fabstatsshow`

Description Use this command to display statistics for the fabric. The following information is displayed:

- Number of times a switch domain ID has been forcibly changed
- Number of E_Port offline transitions
- Number of fabric reconfigurations
- Number of fabric segmentations resulting from any of the following causes:
 - Loopback
 - Incompatibility
 - Overlap
 - Zoning
 - E_Port segment
 - Licensing
 - Disabled E_Port
 - Platform DB
 - Security incompatibility
 - Security violation
 - ECP error
 - Duplicate WWN
 - E_Port isolated
 - MSFR/RD H&T WWN conflict (internal only)
 - Enhanced/Overlapping TI zones (ETIZ) Incompatibility
 - Exchange Switch Capabilities (ESC) detected conflict
 - Encryption conflict - Both ends are not configured for encryption
 - Compression conflict - Both ends are not configured for encryption
 - Encryp/Comp bw availability - Encryption or compression - Configured but lack of bandwidth in the ASIC

For each recorded incident, the command provides the following additional information:

Count

The total number of times the specific event occurred on various ports on the switch.

Port

The number of the port where the latest incident occurred. An arrow (<) next to the port number denotes the type of event that occurred last.

Timestamp

The time when the latest incident occurred.

fabStatsShow

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the fabric statistics:

```
switch:admin> fabstatsshow
```

| Description | Count | Port | Timestamp |
|------------------------------|-------|-------|--------------------------|
| Domain ID forcibly changed: | 0 | | |
| E_Port offline transitions: | 0 | | |
| Reconfigurations: | 2 | 0 | Sat Dec 18 14:29:56 2010 |
| Segmentations due to: | | | |
| Loopback: | 0 | | |
| Incompatibility: | 0 | | |
| Overlap: | 0 | | |
| Zoning: | 0 | | |
| E_Port Segment: | 0 | | |
| Licensing: | 0 | | |
| Disabled E_Port: | 0 | | |
| Platform DB: | 0 | | |
| Sec Incompatibility: | 0 | | |
| Sec Violation: | 0 | | |
| ECP Error: | 0 | | |
| Duplicate WWN: | 0 | | |
| Eport Isolated: | 0 | | |
| AD header conflict: | 0 | | |
| VF AD conflict: | 0 | | |
| MSFR/RD H&T WWN conflict: | 0 | | |
| ETIZ Incompatibility: | 0 | | |
| ESC detected conflict: | 0 | | |
| Encryption conflict: | 0 | | |
| Compression conflict: | 0 | < 349 | Wed Oct 27 15:33:58 2010 |
| Encryp/Comp bw availability: | 0 | | |

'<' - Denotes the type of event that occurred last.

SEE ALSO [fabRetryShow](#)

fanDisable

Disables a fan unit.

Synopsis `fanisable unit`

Description Use this command to disable a nonfaulty fan unit by setting the RPM speed to 0.

Notes This command is not available on nonbladed systems.

Disabling the fan units on a Brocade X6 Director does not change their RPM speed to 0.

Fan units on a Brocade X6 Director take few minutes to change the speed to minimum RPM after being disabled.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

unit
Specifies the number of the fan unit to disable.

Examples To disable a fan unit:

```
switch:admin> fanisable 1
```

```
Fan unit 1 has been disabled
```

SEE ALSO [fanEnable](#), [fanShow](#)

fanEnable

Enables a fan unit.

Synopsis `fanenable unit`

Description Use this command to set a previously disabled fan unit back to the default RPM speed.

Notes This command is not available on nonbladed systems.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

unit

Specify the fan unit number to enable.

Examples To enable a fan unit:

```
switch:admin> fanenable 1
```

```
Fan unit 1 has been enabled
```

SEE ALSO [fanDisable](#), [fanShow](#)

fanShow

Displays fan status and speed.

Synopsis `fanshow`

Description Use this command to display the current status and speed of each fan in the system.

Fan status is displayed as follows:

| | |
|----------------------|---|
| OK | Fan is functioning correctly. |
| absent | Fan is not present. |
| below minimum | Fan is present but rotating too slowly or stopped. |
| above maximum | Fan is rotating too quickly. |
| unknown | Unknown fan unit installed. |
| faulty | Fan has exceeded hardware tolerance and has stopped. In this case, the last known fan speed is displayed. |

The output from this command varies depending on switch type and number of fans present.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display information on the fans in the system:

```
switch:admin> fanshow
Fan #1 is OK, speed is 2721 RPM
Fan #2 is OK, speed is 2657 RPM
Fan #3 is OK, speed is 2700 RPM
```

SEE ALSO [chassisShow](#), [fanDisable](#), [fanEnable](#), [psShow](#)

faPwwn

Manages fabric-assigned port world wide names.

Synopsis

```

fapwwn --enable -port [slot/]port[-port]
fapwwn --enable -ag AG_WWN] -port port
fapwwn --disable -port [slot/]port[-port]
fapwwn --disable -ag AG_WWN -port port
fapwwn --assign [-ag AG_WWN] -port [slot/]port [-v VPWWN]
fapwwn --delete -port [slot/]port[-port]
fapwwn --delete -ag AG_WWN [-port port]
fapwwn --delete -v VWWN
fapwwn --move -srcport source_port -dstport dest_port
fapwwn --move -srcag AG_WWN -srcport source_port
    -dstag AG_WWN -dstport dest_port
fapwwn --show [-port | -ag ] all
fapwwn --show -port [slot/]port[-port]
fapwwn --show -ag AG_WWN [-port]
fapwwn --help

```

Description

Use this command to create and manage fabric-assigned port world wide names (FA-PPWWNs) for Dynamic Fabric Provisioning (DFP). A FA-PWWN is a "virtual" port WWN that will bind to a device port and can be later assigned to a physical device. A FA-PWWN can be either user-generated or automatically assigned by the fabric. The automatically assigned FA-PWWN is created by default when you enable the feature without explicitly providing a VPWWN. With FA-PWWNs assigned to FC ports or Access Gateway (AG) ports, you can provision the ports in advance with zoning configurations, access control, or Ports on Demand assignments, and you can easily move servers across ports or Access Gateways by way of reassigning the FA-PWWN to another port.

The **faPwwn** command supports the following management tasks:

- Bind an automatically assigned or a user-assigned FA-PWWN to a switch port.
- Override an automatically assigned FA-PWWN with a user-assigned FA-PWWN.
- Bind an AG port with an automatically assigned or a user-assigned FA-PWWN.
- Delete any existing FA-PWWN bindings.
- Move a FA-PWWN from one port to another port.
- Move a FA-PWWN assigned to an AG port to another AG.
- Display information about configured FA-PWWN bindings.

A single port can be assigned up to two WWNs, one assigned automatically and one assigned by the user. Only one FA-PWWN can be active at any given time. The user-assigned FA-PWWN takes precedence over the automatically assigned FA-PWWN. This means, the switch will bind the user-assigned FA-PWWN to the port if both a user-assigned and an automatically assigned FA-PWWN are available.

FA-PWWN configurations are saved persistently and are preserved even if the feature is disabled on a port, or if you move the port to a different logical switch. The same configuration will be reapplied on the port once the feature is enabled or the port is moved back to the switch where you initially created and saved the configuration.

When issued for switch ports, the **--show** option displays the following information:

| | |
|--------------|---------------------------|
| Port | Port number |
| PPWNN | Real PWWN of the device |
| VPWNN | Currently active FA-PWWN. |

The following information is displayed for Access Gateway ports:

| | |
|-------------------------|---|
| AG Port | Access Gateway node WWN followed by a slash and port number |
| Port | The FC switch port number |
| Device Port WWN | The real device port WWN displays only if the device is logged into the Access Gateway. |
| Virtual Port WWN | The FA-PWWN (user-assigned or automatically assigned) created on the FC switch for the Access Gateway port. |

The following information is displayed for both AG and FC switch ports:

| | |
|----------------|---|
| PID | Port ID. |
| Enable | Displays "yes" if the configuration is enabled; displays "No" if the configuration is disabled. |
| MapType | Displays the type of mapping: Access Gateway (AG) or switch port (port), user-assigned (user) or automatically assigned (auto). |

Notes This command is supported only on the Brocade DCX 8510-4 and DCX 8510-8, and the Brocade 6510. The AG port must be configured on the switch to which the AG is connected. However, both AG switches are FA-PWWN-aware, which means that they can connect to switches with configured FA-PWWN ports. Refer to the Release Notes for other restrictions that may apply.

You cannot enable the FA-PWWN feature if **portSwap** is enabled on the ports.

The operands, **--assign**, **--delete**, **--enable**, and **--disable** can be executed only on a disabled port (switch ports and AG ports). You must re-enable the ports to come online with the new configuration.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands this command has the following operands:

| | |
|------------------|--|
| --enable | Enables the FA-PWWN feature on the specified switch ports or AG ports. If a FA-PWWN is not available for the ports, this command automatically assigns FA-PWWNs to the ports. The switch ports must be disabled using the portDisable command before executing this command. |
| --disable | Disables the FA-PWWN feature on the specified switch ports or AG ports. This command disables the binding of any existing FA-PWWNs to the specified ports, it does not delete the port configurations themselves. When you re-enable the feature, the persistently saved configurations are re-applied to the ports. |
| --assign | Assigns a FA-PWWN to a specified port. You can specify only a single port with this command (not a port range) and optionally provide a VPWWN. The switch ports must be disabled using the portDisable command before executing this command. When executed without the -v option, this command automatically generates a FA-PWWN and assigns it to the port. When executed with the -v option, and a user-assigned FA-PWWN already exists, this command fails with an appropriate message. If an automatically assigned FA-PWWN exists, the user-generated FA-PWWN overrides the automatically generated FA-PWWN. When you assign a FA-PWWN to an AG port, you must specify the AG node WWN to uniquely identify the AG. The new FA-PWWN will take effect upon the next login session. |

--delete

Deletes the active FA-PWWN from the specified ports. If a port is bound to a user-assigned FA-PWWN, this command deletes the currently active FA-PWWN and activates the automatically assigned FA-PWWN. If the active FA-PWWN is automatically assigned and no user-assigned FA-PWWN exists, this command deletes the FA-PWWN. When deleting VPWWNs from an AG, you must specify the AG node WWN and the **-port** operand is optional; if omitted, all active FA-PWWNs are deleted from the AG. To delete a single FA-PWWN, use the delete command with the **-v** option.

If no active FA-PWWNs exist on the ports, the FA-PWWN feature is disabled. Use the **-enable** command to re-enable the feature.

--show

Displays the FA-PWWN configuration for the specified ports or for all ports in the logical switch. When displaying information for AG ports, you must specify the AG node WWN.

-port[slot]port[-port]

Specifies the ports to be configured or displayed on the switch or on the AG, preceded by a slot number on bladed systems. The port list can include a single port or a port range, for example, 3-5 or 1/3-5. Port ranges are supported only with automatically assigned FA-PWWNs, and only on switch ports. You cannot use a port range on AG ports.

-ag AG_WWN

Specifies the Access Gateway node WWN.

-v VPWWN

Specifies a virtual PWWN to be assigned to the port. This operation deactivates any existing FA-PWWN on the port.

all

Specifies all ports on the logical switch. This operand is valid only with the **--show** option.

--move

Moves an active FA-PWWN from a source to the specified destination port. Use this command to move a server across switch ports. If the source port has both a user-assigned and an automatically assigned FA-PWWN, the user-assigned FA-PWWN will be moved to the destination port as a user-assigned FA-PWWN and replace any automatically assigned FA-PWWN that may be active on that port. If the source port has only an auto-assigned FA-PWWN, the auto-assigned FA-PWWN will be moved to the destination port and the FA-PWWN feature will be disabled on the source port, because there is no FA-PWWN left on that port. When moving PWWNs between Access Gateways, you must specify the source AG node WWN and the destination AG node WWN in addition to the ports.

The following operands are supported with the **--move** option:

-srcag AG_WWN

Specifies the source AG by its node WWN. This operand is required only when moving ports between Access Gateways.

-dstag AG_WWN

Specifies the destination AG by its node WWN. This operand is required only when moving ports between Access Gateways.

-srcport *source_port*
Specifies the port from which to move the PWWN.

-dstport *dest_port*
Specifies the destination port for the moved PWWN.

--help
Displays the command usage.

Examples

To enable the FA-PWWN feature on a switch port:

```
switch:admin> portdisable 10
switch:admin> fapwwn --enable -port 10
```

To disable the FA-PWWN feature on the switch port:

```
switch:admin> fapwwn --disable -port 10
```

To move the active FA-PWWN from a source port to another port:

```
switch:admin> fapwwn --move -srcport 2 -dstport 3
```

To move the active FA-PWWN from a source AG port to a destination AG port:

```
switch:admin> fapwwn --move -srcag 11:11:11:11:11:11 \
    -srcport 12 -dstag 22:22:22:22:22:22 -dstport 14
```

To assign a user-assigned FA-PWWN to a switch port:

```
switch:admin> portdisable 1/10
switch:admin> fapwwn --assign -port 1/10 -v 50:60:70:71:72:73:81:81
```

To assign an automatically assigned FA-PWWN to a switch port:

```
switch:admin> portdisable 1/10
switch:admin> fapwwn --assign -port 1/10
```

To assign a fabric-assigned FA-PWWN to an AG port identified by the AG node WWNN and port number:

```
switch:admin> portdisable 0
switch:admin> fapwwn --assign -ag 12:34:56:78:90:12:23:45 -port 0
```

To assign an user-assigned FA-PWWN to an AG port:

```
switch:admin> portdisable 0
switch:admin> fapwwn --assign -ag 12:34:56:78:90:12:23:45 -port 0 \
    -v 11:22:33:44:55:66:77:88
```

To delete the active FA-PWWN of a switch port:

```
switch:admin> fapwwn --delete -port 10
```

To delete the active FA-PWWN of an AG port:

```
switch:admin> fapwwn --delete -ag 12:34:56:78:90:12:34:56 -port 10
```

To display the active FA-PWWN for a single FC port (the real device PWWN is hidden):

```
switch:admin> fapwwn --show -port 10
-----
Port          PPWWN                      VPWWN                      PID Enable MapType
-----
10  ---:---:---:---:---:---:---  52:00:10:00:00:0f:50:45  --  Yes  Port/Auto
```

To display the active FA-PWWNs for all FC switch ports:

```
switch:admin> fapwwn --show -port all
-----
Port          PPWWN                      VPWWN                      PID Enable MapType
-----
0  ---:---:---:---:---:---:---  52:00:10:00:00:0f:50:30  10101  Yes  Port/Auto
1  ---:---:---:---:---:---:---  11:22:33:44:33:22:11:22  --  Yes  Port/User
                               52:00:10:00:00:0f:50:44
10 ---:---:---:---:---:---:---  52:00:10:00:00:0f:50:45  --  Yes  Port/Auto
```

To display the active FA-PWWN of an AG port :

```
switch:admin> fapwwn --show -ag 10:00:00:05:1e:d7:3d:dc -port 8
-----
AG Port          Port          Device Port WWN  \
-----
10:00:00:05:1e:d7:3d:dc/8  20  20:08:00:05:1e:d7:2b:74  \
-----
Virtual Port WWN          PID  Enable  MapType
-----
11:22:33:44:55:66:77:88  11403  Yes  AG/User
52:00:10:00:00:0f:50:32
(output split)
```

To display the active FA-PWWNs of all configured AG ports:

```
switch:admin> fapwwn --show -ag all
-----
AG Port          Port          Device Port WWN  \
-----
10:00:00:05:1e:65:8a:d5/16  --  ---:---:---:---:---:---:---  \
10:00:00:05:1e:d7:3d:dc/8  20  20:08:00:05:1e:d7:2b:74  \
-----
10:00:00:05:1e:d7:3d:dc/9  20  20:09:00:05:1e:d7:2b:73  \
10:00:00:05:1e:d7:3d:dc/16  --  ---:---:---:---:---:---:---  \
-----
Virtual Port WWN          PID  Enable  MapType
-----
52:00:10:00:00:0f:50:30  --  Yes  AG/Auto
11:22:33:44:55:66:77:88  11403  Yes  AG/User
52:00:10:00:00:0f:50:32
2:00:10:00:00:0f:50:33  11404  Yes  AG/Auto
52:00:10:00:00:0f:50:38  --  Yes  AG/Auto
(output split)
```

To display the active FA-PWWNs of all ports associated with a single AG (in the following example, one VPWWN is not unassigned):

```
switch:admin> fapwwn --show -ag 10:00:00:05:1e:d7:3d:dc
-----
      AG Port                Port      Device Port WWN      \
-----
10:00:00:05:1e:d7:3d:dc/16  --      --:--:--:--:--:--:-- \
10:00:00:05:1e:d7:3d:dc/9   20      20:09:00:05:1e:d7:2b:73 \
10:00:00:05:1e:d7:3d:dc/8   20      20:08:00:05:1e:d7:2b:74 \
\
10:00:00:05:1e:d7:3d:dc/16  --      --:--:--:--:--:--:-- \
-----
      Virtual Port WWN      PID      Enable  MapType
-----
52:00:10:00:00:0f:50:38    --      Yes     AG/Auto
20:00:10:00:00:0f:50:33    11404   Yes     AG/Auto
11:22:33:44:55:66:77:88    11403   Yes     AG/User
52:00:10:00:00:0f:50:32
(output split)
```

SEE ALSO None

fastBoot

Reboots the Control Processor (CP), bypassing the Power-On Self-Test (POST).

Synopsis `fastboot`

Description Use this command to perform a "cold reboot" (power off/restart) of the CP bypassing POST when the system comes back up. Bypassing POST can reduce boot time significantly. If POST was previously disabled using the `diagDisablePost` command, then `fastBoot` is the same as `reBoot`.

The `fastBoot` operation is disruptive, and the command prompts for confirmation before executing. When you reboot a switch connected to a fabric, all traffic to and from that switch stops. All Fibre Channel ports on that switch including E_Ports become inactive until the switch comes back online.

The behavior of this command varies according to platform type:

- When issued on a standalone (single-processor) switch, this command performs a cold reboot of the switch.
- When issued on an enterprise-class platform (Brocade DCX, DCX-4S, or 48000) with two CPs (active and standby), the following rules apply:
 - When the Standby CP reboots, it goes down and there is no failover because there is no traffic on that switch. When the Standby CP comes up again, it is temporarily no longer in sync with the Active CP.
 - When the Active CP reboots, it fails over to the Standby CP. The Standby CP becomes the new Active CP and traffic is disrupted.
 - When HA is disabled and `fastBoot` is issued on the Active CP, both the Active and Standby CPs reboot with the original mastership retained. The original Active CP remains the Active CP after the reboot, and the original Standby CP remains the Standby CP. After the reboot, HA is enabled.
 - When HA is disabled and `fastBoot` is issued on the Standby CP, the Standby CP reboots without prompting. It boots up with the default switch only, even if the Active CP has multiple logical switches configured. After the Standby CP boots up, HA is still disabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To reboot a standalone (single-processor) switch without executing POST on startup:

```
switch:admin> fastboot

Warning: This command would cause the switch to reboot
and result in traffic disruption.
Are you sure you want to reboot the switch [y/n]? y

Broadcast message from root (pts/0) Sun Feb 28 19:49:45 2010...

The system is going down for reboot NOW !!
```

fastBoot

To reboot a DCX without executing POST on startup (in the example, HA is enabled):

```
switch:admin> fastboot
Warning: This command is being run on a control processor (CP)
based system and will cause the active CP to reboot.
Are you sure you want to reboot the active CP [y/n]? y

Broadcast message from root (pts/0) Sun Feb 28 19:49:45 2010...

The system is going down for reboot NOW !!
```

To reboot a DCX without executing POST on startup (in the example, HA is disabled):

```
switch:admin> fastboot
This command is being run on a control processor (CP)
based system. Because HA is disabled, it will cause both
active CP and the standby CP to reboot. After reboot, the
HA will be enabled.

Do you want to continue [y/n] y

Broadcast message from root (pts/0) Sun Feb 28 19:49:45 2010...

The system is going down for reboot NOW !!
```

SEE ALSO [diagDisablePost](#), [diagEnablePost](#), [reboot](#), [haDisable](#), [haEnable](#), [haFailover](#)

fcipHelp

Displays FCIP command information.

Synopsis `fciphelp`

Description Use this command to display a listing of Fibre Channel over IP (FCIP) commands with short descriptions for each command. FCIP commands require an FCIP license.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display FCIP command help information:

```
switch:admin> fciphelp
bladecfggemode      Configure 1GigE/10GigE port configuration for blade
extncfg             Configure blade/switch extension product features
fciphelp            Print FCIP help info
licenseslotcfg     Configure licenses on a slot basis
portcfg            Create/Delete a new ip interface/route/arp entry on
                   the GigE port
portcfgge          Configure GigE port features
portcfggemediatype Configure media type for GigE port
portcmd            Execute commands (ping etc) on the GigE port
portshow           Show configured ip interfaces/routes/arp entries on
                   the GigE Port
```

SEE ALSO [portCfg](#), [portCmd](#), [portShow](#)

fcipLedTest

Exercises the GbE port LEDs on the Brocade 7840, FX8-24, and SX6 extension blade.

Synopsis `fcipledtest [--slot slot | all]`

Description Use this command to exercise the GbE port LEDs on the Brocade Brocade 7840, FX8-24, and SX6 extension blade. The behavior of this command varies according to platform:

- On Brocade FX8-24, this test cycles through the port LEDs by lighting GREEN and then flashing GREEN on all ports for 3 seconds. As the test continues the ports turn AMBER and then flashing AMBER for 3 seconds. The LEDs turn off when the test has finished.
- On Brocade 7840, this test cycles all the port LEDs by lighting GREEN on all ports for 3 seconds. The switch will reboot once the test is complete.
- On Brocade SX6, this test cycles all the port LEDs by lighting GREEN on all ports for 3 seconds. As the test continues the ports turn AMBER for 3 seconds. The LEDs turn off when the test has finished. The blade will reboot once the test is complete.

This diagnostic cannot be run on an operational switch. You must disable the switch using the **chassisDisable** command before you can run this test. After the command completes, the ATTN LEDs flash amber, indicating that the command has finished and exited. Enable the switch using the **chassisEnable** command to set the ATTN LEDs back to black.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--slot *slot*

Tests all GbE ports in the specified slot. This operand is optional. This operand is valid only on chassis-based platforms.

all

Tests all GbE ports in the chassis. This operand is optional. This operand is valid only on chassis-based platforms.

Examples To test the LEDs on slot 4 of Brocade DCX 8510 with an FX8-24 blade.

```
switch:admin> fcipledtest --slot 4
PASSED.
```

SEE ALSO [bladeDisable](#), [bladeEnable](#), [chassisDisable](#), [chassisEnable](#), [extnCfgr](#), [portLedTest](#)

fcipPathTest

Tests the data path of the FCIP complex.

Synopsis `fcipathtest`
 `[--slot slot]`
 `[-lb_mode mode]`
 `[-nframes count]`

Description Use this command to verify the data paths in the FCIP complex. All data path modes run tests by comparing Fibre Channel frames or data packets transmitted from and received by the network processor due to the designated loopback.

This command is supported only on the Brocade 7840, FX8-24, SX6 platforms.

Executing this command causes the switch or blade to reboot.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------------------------|---|
| <code>--slot <i>slot</i></code> | Specifies the slot number on which the diagnostic operates. The default is 0 and operates on fixed-port-count products. |
| <code>-lb_mode <i>mode</i></code> | Specifies the loopback mode for the test. By default, this test uses the External (SERDES) loopback mode 2 which is the only loopback mode supported on Brocade 7840 and SX6 platforms. Other valid values in addition to mode 2 for FX8-24 platforms are as follows: |
| 1 | Port loopback (requires loopback plugs) |
| 2 | External (SERDES) loopback |
| 5 | Internal (parallel) loopback |
| 7 | Backend bypass and port loopback (requires loopback plugs) |
| 8 | Backend bypass and SERDES loopback |
| <code>-nframes <i>count</i></code> | Specifies the number of frames to send. The test progresses until the specified number of frames is transmitted on each port. The default value is 100 for FX8-24 platforms and 4 for Brocade 7840 platform. |

fcipPathTest

Diagnostics When it detects failures, the test may report one or more of the following error messages:

PATH_TEST_ERR
PATH_TEST_CHIP_INIT_ERR
PATH_TEST_IMAGE_ERR
PATH_TEST_TIMEOUT_ERR
PATH_TEST_HEARTBEAT_ERR
PATH_TEST_INVALID_RESULT
PATH_TEST_GE_PORT_ENABLE_ERR
PATH_TEST_GE_PORT_DISABLE_ERR

Examples To run the test on slot 2:

```
switch:admin> fcippathtest --slot 2
Running fcippathtest .....
Test Complete: fcippathtest Pass 10 of 10
Duration 0 hr, 1 min & 50 sec (0:1:50:942).
passed.
```

SEE ALSO [bladeDisable](#), [bladeEnable](#), [chassisDisable](#), [chassisEnable](#), [extnCfg](#)

fcPing

Sends a Fibre Channel Extended Link Service (ELS) Echo request to a pair of ports or to a single destination, or executes a SuperPing.

Synopsis

```
fcping
  [--number frames]
  [--length size]
  [--interval wait]
  [--pattern pattern]
  [--bypasszone]
  [--quiet]
  [source] destination
```

```
fcping --allpaths
  [-printis!]
  [-maxtries M]
  [-covcount M]
  [-delay D]
  [-framelength len]
  [-errstats]
  [-vc]
  destination
```

```
fcping --help
```

Description

Use this command to send a Fibre Channel ELS Echo request to a pair of ports (a source and a destination), to a single device, or to execute a SuperPing that exercises all interswitch links (ISLs) and internal links in different paths that route to the destination device.

- When you use **fcPing** with a source and a destination, the command performs a zoning check between the two ports. In addition, two Fibre Channel ELS requests are generated. The first ELS request is from the domain controller to the source port identifier. The second ELS request is from the domain controller to the destination port identifier. The ELS Echo request elicits an ELS Echo response from a port identifier in the fabric and is useful for validating link connectivity.

The source and destination port identifiers can be specified as a 24-bit Fibre Channel port identifier (PISD), a port World Wide Name, or a node World Wide Name. The two port identifiers are then used to determine if the identifiers are zoned together.
- When you use **fcPing** to probe a single destination, an ELS Echo is sent to the specified destination and a response obtained. The destination can be a switch WWN, a domain ID, or a switch domain controller ID. No zoning check is performed when a single device is probed.
- When you use **fcPing** with the **--allpaths** option, the command exercises a "SuperPing." A SuperPing exercises all ISLs and the internal links included in the least-cost paths that route to the destination. It collects statistical data of all the covered paths and their port and provides optional parameters to selectively display the data. SuperPing takes only one argument, the destination port identifier. To execute a SuperPing for two destinations, you must issue the **fcping --allpaths** command separately for each destination.

SuperPing facilitates troubleshooting of links that experience problems. When an echo frame is dropped, all the ISLs and internal links potentially traversed by this frame are marked as failures. If a fabric topology is considered fully redundant, that is, at each hop there are multiple paths to reach a destination, a high percentage of errors are recorded on the link that experiences errors.

Logical Fabrics: When executed in a Logical Fabric from a switch to a destination device connected through the base fabric, SuperPing exercises all paths in the base fabric along with the ISLs in the logical fabric. The path output indicates the LISLs and the base switch. Refer to the Examples section for an illustration.

Notes The ELS Echo may not be supported on all devices. In such cases, the response could be either an ELS reject or a request timeout.

By default, **fcPing** sends five ELS Echo requests to each port. When a device does not respond to the ELS Echo request, further debugging may be needed to determine, whether the device does not support ELS Echo, or whether the request is rejected for some other reason. Do not assume that the device is not connected.

If a fabric reconfiguration occurs while SuperPing is in progress, the command reports an error message. Exit the command and rerun the test after the fabric becomes stable again.

SuperPing and the **fcPing** command are not supported for the simulated devices that exist on SIM ports. Use the **flow** command to enable the SIM ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

destination

Specifies the destination as follows.

- When using **fcPing** between a source and a destination, specify the destination as port WWN or a node WWN.
- When using **fcPing** to ping a single device, specify the destination as a PID, a switch WWN, a domain ID, or a switch domain controller ID.
- When using **fcPing** with the **--allpaths** option, specify the destination as a PID, a switch WWN, or a domain ID.

source

Specifies the source port ID, port WWN, or node WWN. This operand is optional; it is not valid with the SuperPing command (**--allpaths**).

The following operands are valid only when **fcPing** is executed without the **--allpaths** option (legacy mode):

--number frames

Specifies the number of ELS Echo requests to send. The default value is 5.

--length size

Specifies the frame size of the requests in bytes. The default value is 0. Without data, the Fibre Channel Echo request frame size is 12 bytes. The total byte count includes four bytes from the Echo request header and eight bytes from the timestamp. The maximum allowed value is 2,036 bytes. The length must be word-aligned.

--interval wait

Specifies the interval, in seconds, between successive ELS Echo requests. The default value is 0 seconds.

- pattern *pattern***
Specifies up to 16 "pad" bytes, which are used to fill out the request frame payload sent. This is useful for diagnosing data-dependent problems in the fabric link. The pattern bytes are specified as hexadecimal characters. For example, **--pattern ff** fills the request frame with instances of the number 1. If a non-byte-aligned pattern is specified, the upper nibble of the last pattern byte is filled with zeros. For example, **--pattern 123** fills the payload with a pattern of 0x1203.
- bypasszone**
Bypasses the zone check.
- quiet**
Suppresses the diagnostic output. Only zoning information, if applicable, and the summary line are displayed.

The following operands are valid only when **fcPing** is executed to perform a SuperPing:

- allpaths [*args*] *destination***
Executes a SuperPing that covers all available least-cost paths to the specified destination. The number of actual paths covered depends on two other parameters that you can optionally specify. When you issue **fcping --allpaths** for a destination without any other options, SuperPing covers all ISLs in the routes between source to destination but does not exercise all possible combinations of end-to-end paths. This operand is required when executing **fcPing** as SuperPing.

The following operands are optional and valid only with the **--allpaths** option:

- printisl**
Displays statistical data for each ISL and internal port along the paths traversed by SuperPing. This information displays in addition to the path display.
- covcount *N***
Specifies the minimum number of times each ISL is exercised by the SuperPing command. The command sends *N* frames and checks if each ISL is exercised at least *N* times. When the condition is met, superPing exits and prints the statistics. The default value is 5.
- maxtries *M***
Specifies the maximum number of frames to be sent before SuperPing exits. If both **-maxtries** and **-covcount** are specified, SuperPing checks the ISL coverage and keeps resending frames until the minimum coverage condition is met or until the maximum number of echo frames specified in **maxtries** has been sent. For example, assuming a coverage count of 100 and a **Maxtries** value of 300, SuperPing will send 100 frames at a time and checks if each ISL is covered at least 100 times. If not, SuperPing will keep sending 100 frames at a time to check for coverage up to 3 times for a maximum of 300 frames (3*100 = 300) on each egress port. If this value is set too low in relation to the specified coverage count, not all ISLs may be covered. The default value is 100.
- delay *D***
Includes a delay of *D* milliseconds between each echo frame sent.
- framelength *len***
Specifies the size of the data to send.

- errstats** Collects error statistics of each user port that is part of ISLs covered through SuperPing to reach destination domain.
- vc** Displays the ISL and internal port statistics per VC.
- help** Displays the command usage.

Examples

To display one device that accepts the request and another device that rejects the request:

```
switch:admin> fcping 10:00:00:00:c9:29:0e:c4 21:00:00:20:37:25:ad:05
Source:          10:00:00:00:c9:29:0e:c4
Destination:     21:00:00:20:37:25:ad:05
Zone Check:     Not Zoned

Pinging 10:00:00:00:c9:29:0e:c4 [0x20800] with 12 bytes of data:
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1162 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1013 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1442 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1052 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec

Pinging 21:00:00:20:37:25:ad:05 [0x211e8] with 12 bytes of data:
Request rejected
Request rejected
Request rejected
Request rejected
Request rejected
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 0/0/0 usec
```

To display one device that accepts the request and another device that does not respond to the request:

```
switch:admin> fcping 0x020800 22:00:00:04:cf:75:63:85
Source:          0x020800
Destination:     22:00:00:04:cf:75:63:85
Zone Check:     Zoned

Pinging 0x020800 with 12 bytes of data:
received reply from 0x020800: 12 bytes time:1159 usec
received reply from 0x020800: 12 bytes time:1006 usec
received reply from 0x020800: 12 bytes time:1008 usec
received reply from 0x020800: 12 bytes time:1038 usec
received reply from 0x020800: 12 bytes time:1010 usec
5 frames sent, 5 frames received, 0 frames rejected,0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec

Pinging 22:00:00:04:cf:75:63:85 [0x0217d9] with 12 bytes of data:
Request timed out
Request timed out
Request timed out
Request timed out
Request timed out
5 frames sent, 0 frames received, 0 frames rejected,5 frames timeout
Round-trip min/avg/max = 0/0/0 usec
```

To use **fcping** with a single destination:

```
switch:admin> fcping 20:03:00:05:33:7e:97:e9
Pinging 20:03:00:05:33:7e:97:e9 [0x022300] with 12 bytes of data:
Request rejected
Request rejected by 0x022300: Command not supported: time: 1159 usec
Request rejected by 0x022300: Command not supported: time: 1006 usec
Request rejected by 0x022300: Command not supported: time: 1008 usec
Request rejected by 0x022300: Command not supported: time: 1038 usec
Request rejected by 0x022300: Command not supported: time: 1010 usec
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec
```

To use **fcPing** with a single destination (in the example, the destination is a switch WWN):

```
switch:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Name
-----
6: fffc06 10:00:00:05:1e:34:2b:66 10.202.90.201 0.0.0.0 "mps_daz_1"
55:fffc37 10:00:00:05:1e:34:01:f5 10.202.90.226 0.0.0.0 pulsar055"

switch:admin> fcping 10:00:00:05:1e:34:2b:66
Destination:      10:00:00:05:1e:34:2b:66

Pinging 20:00:00:05:1e:34:2b:66 [fffc06] with 12 bytes of data:
received reply from 10:00:00:05:1e:34:2b:66 : 12 bytes time:1162 usec
received reply from 10:00:00:05:1e:34:2b:66 : 12 bytes time:1013 usec
received reply from 10:00:00:05:1e:34:2b:66 : 12 bytes time:1442 usec
received reply from 10:00:00:05:1e:34:2b:66 : 12 bytes time:1052 usec
received reply from 10:00:00:05:1e:34:2b:66 : 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec
```

To use **fcPing** with a single destination (in the example, the destination is a device node WWN):

```
switch:admin> nsshow
{
Type Pid      COS      PortName      NodeName      TTL(sec)
N 370500;3;20:07:00:05:1e:35:10:7f;10:00:00:05:1e:35:10:7f; na
Fabric Port Name: 20:05:00:05:1e:34:01:f5
Permanent Port Name: 20:07:00:05:1e:35:10:7f
Port Index: 5
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: NO
N 370501; 3;10:00:00:00:c9:3f:7c:b8;20:00:00:00:c9:3f:7c:b8; na
FC4s: FCP
NodeSymb: [44] "Emulex LP1050 FV1.81A1 DV5-5.20A9 DELL1750-3"
Fabric Port Name: 20:05:00:05:1e:34:01:f5
Permanent Port Name: 20:07:00:05:1e:35:10:7f
Port Index: 5
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: NO
The Local Name Server has 2 entries }
```

```

switch:admin> fcping 20:00:00:00:c9:3f:7c:b8
Destination:    20:00:00:00:c9:3f:7c:b8

Pinging 20:00:00:00:c9:3f:7c:b8 [0x370501] with 12 bytes of data:
received reply from 20:00:00:00:c9:3f:7c:b8:12 bytes time:825 usec
received reply from 20:00:00:00:c9:3f:7c:b8:12 bytes time:713 usec
received reply from 20:00:00:00:c9:3f:7c:b8:12 bytes time:714 usec
received reply from 20:00:00:00:c9:3f:7c:b8:12 bytes time:741 usec
received reply from 20:00:00:00:c9:3f:7c:b8:12 bytes time:880 usec
5 frames sent,5 frames received,0 frames rejected,0 frames timeout
Round-trip min/avg/max = 713/774/880 usec

```

To execute a SuperPing testing all ISLs to a specified destination (in the example, two paths are tested, and each hop is displayed in Domain/Index format):

```

switch:admin> fcping --allpaths 165

Pinging(size:12 bytes) destination domain 165 through all paths

PATH SWITCH1-->    SWITCH2-->    SWITCH3-->    \
-----
1.(3/EMB,3/205)[128] (207/25,207/42)[128] (101/3,101/16)[128]\
2.(3/EMB,3/204)[128] (207/27,207/42)[128] (101/3,101/16)[128]\

SWITCH4          STATUS
-----
(165/99,165/0)[128] SUCCESS
2(165/99,165/0)[128] SUCCESS

```

To execute a SuperPing in a logical fabric (in the example, domains 10 and 40 in FID 1 are connected through the base fabric (FID 2)):

```

switch:admin: fcping --allpaths 40
Pinging(size:12 bytes) destination domain 30 through all paths

PATH SWITCH1-->    SWITCH2-->
-----
1.(10/EMB,10/4)[128] (20/5,20/EMB ) [128] (1/EMB, 1/6)[2] \
Successfully completed superping for all paths

SWITCH3-->          SWITCH4          STATUS
-----
(2/7,2/EMB)[2] (30/EMB,30/8)[128] (40/9, 40/EMB)[128] SUCCESS
Successfully completed superping for all paths

```

To execute a SuperPing and print statistical coverage of each ISL and internal port along the potential paths (in the example, a few errors are recorded on the ISLs 3/205->2/25, 3/204->2/27, 2/42->101/3, and 2/1->101/8, but the internal port analysis shows that errors are recorded on the internal port 0/284 in domain 2, which is the potential faulty link):

```

switch:admin> fcping --allpaths -printisl 101
Pinging(size:12 bytes) destination domain 101 through all paths

PATH SWITCH1-->    SWITCH2-->    \
-----
1. ( 3/EMB, 3/123)[128] (165/96 ,165/99 ) [128] \

```

```

2. ( 3/EMB, 3/205)[128] ( 2/25 , 2/1 ) [128] \
3. ( 3/EMB, 3/205)[128] ( 2/25 , 2/42 ) [128] \
4. ( 3/EMB, 3/204)[128] ( 2/27 , 2/42 ) [128] \
5. ( 3/EMB, 3/204)[128] ( 2/27 , 2/1 ) [128] \

```

```

SWITCH3                STATUS
-----
(101/16 ,101/EMB)[128] SUCCESS
(101/8 ,101/EMB)[128]  FAILED
(101/3 ,101/EMB)[128]  FAILED
4(101/3 ,101/EMB)[128] FAILED
 101/8 ,101/EMB)[128]  FAILED

```

Completed superping for all paths. Error found in few paths

ISL COVERAGE

```

-----
SNO                ISL                STATUS
-----
1 ( 3/123[128]-->165/96[128] )          SUCCESS(5/5)
2 ( 3/205[128]--> 2/25[128] )          FAILURE(7/50)
3 ( 3/204[128]--> 2/27[128] )          FAILURE(11/50)
4 (165/99[128] -->101/16[128] )          SUCCESS(5/5)
6 ( 2/42[128] -->101/3[128] )          FAILURE(10/67)
7 ( 2/1[128] -->101/8[128] )           FAILURE(8/33)

```

INTERNAL PORT COVERAGE

```

-----
SNO    DOMAIN  INTRNL_PORT  STATUS
-----
1      2       0/272        SUCCESS(40/40)
2      2       0/276        SUCCESS(44/44)
3      2       0/280        SUCCESS(30/30)
4      2       0/284        FAILURE(20/20)

```

To execute a SuperPing with a coverage count of 1000 and a **maxtries** value of 5000 (in the example, the ISL (3/204->204/27), could not be covered 1000 times):

```

switch:admin> fcping --allpaths -covcount 1000 \
               -maxtries 5000 -printisl 165

Pinging(size:12 bytes) destination domain 165 through all paths
|
PATH SWITCH1--> SWITCH2--> SWITCH3--> \
-----
1.(3/EMB,3/205)[128] (207/25,207/42)[128] (101/3,101/16)[128]\
2.(3/EMB,3/204)[128] (207/27,207/42)[128] (101/3,101/16)[128]\

SWITCH4                STATUS
-----
(165/99,165/0)[128] SUCCESS

(165/99,165/0)[128] SUCCESS

Successfully completed superping for all paths

```

ISL COVERAGE

| SNO | ISL | STATUS |
|-----|---------------------|--------------------|
| 1 | (3/205-->207/25) | SUCCESS(4025/4025) |
| 2 | (3/204-->207/27) | SUCCESS(5/5) |
| 3 | (207/42 -->101/3) | SUCCESS(4030/4030) |
| 4 | (101/16 -->165/99) | SUCCESS(4030/4030) |

INTERNAL PORT COVERAGE

| SNO | DOMAIN | INTRNL_PORT | STATUS |
|-----|--------|-------------|--------|
|-----|--------|-------------|--------|

SEE ALSO None

fcpLogClear

Clears the FCPD debug information log.

Synopsis `fcplogclear`

Description Use this command to clear the debug information logged by the Fibre Channel Protocol daemon (FCPD).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To clear the FCPD debug information log:

```
switch:admin> fcplogclear
```

SEE ALSO [fcpLogDisable](#), [fcpLogEnable](#), [fcpLogShow](#)

fcpLogDisable

Disables the FCPD debug information log.

Synopsis `fcplogdisable`

Description Use this command to disable the logging of debug information by the Fibre Channel Protocol daemon (FCPD).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the FCPD debug information log:

```
switch:admin> fcplogdisable
```

SEE ALSO [fcpLogClear](#), [fcpLogEnable](#), [fcpLogShow](#)

fcpLogEnable

Enables the FCPD debug information log.

Synopsis `fcplogenable`

Description Use this command to enable Fibre Channel Protocol daemon (FCPD) logging. Debug information logging is enabled by default.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable the FCPD debug information log:

```
switch:admin> fcplogenable
```

SEE ALSO [fcpLogClear](#), [fcpLogDisable](#), [fcpLogShow](#)

fcpLogShow

Displays the FCPD debug information log.

Synopsis `fcplogshow`

Description Use this command to display the debug information logged at various stages during the Fibre Channel Protocol daemon (FCPD) device probing.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the FCPD debug information log:

```
switch:admin> fcplogshow
Switch 0; Sat Dec 18 14:22:05 2010
14:22:05.799269 Flush or Probe rcvd, port 15, \
  Line: 16777862, scn_type: 2
14:22:05.799279 Probe flush, port:15, \
  Line:16780554, flag:0
14:22:17.469561 Flush or Probe rcvd, port 15, \
  Line: 16777862, scn_type: 1
14:22:17.469567 Probe Msg Rcvd, port: 15, Line: 16777899, \
  prev_state: 0, msg_scn_arg1: 0
14:22:17.470618 fcpStartProbe, port: 15, line: 16778347, \
  async_state: 1, origin: 0,probe_state: 0, opt_code: 0
14:22:17.471052 fcpStartProbe, port: 15, line: 16778413, \
  async_state: 1073741824, origin: 0,probe_state: 0, opt_code: 0
14:22:17.471057 Probe flush, port:15, Line:16780554, flag:1
(output truncated)
```

SEE ALSO [fcpLogClear](#), [fcpLogDisable](#), [fcpLogEnable](#)

fcpProbeShow

Displays the Fibre Channel Protocol (FCP) probe information.

Synopsis `fcpprobeshow [slot/]port`

Description Use this command to display the Fibre Channel Protocol daemon (FCPD) device probing information for the devices attached to the specified F_Port or FL_Port. This information includes the number of successful logins and SCSI INQUIRY commands sent over this port and a list of the attached devices.

This command includes probing information for NPIV devices. In addition, this command displays the list of devices connected to a port and the number of successful PLOGI, PRLI, INQUIRies, and current probe state. For F_Ports and NPIV ports, when the F_Port Device Update Mode is ON, the probed device details are updated in the "npiv update map" field. When the F_Port Device Update Mode is OFF, the probed device details are updated in the "update map" field. Use the **configure** command to configure the F_Port Device Update Mode.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the port to display, followed by a slash (/).

port

Specify the port number to display, relative to its slot for bladed systems. Use **switchShow** to list valid ports. This operand is required.

Examples To display the FCP probe information:

```
switch:admin> fcpprobeshow 31

port 0 is FV-Port and it is online.
nodes probed:          3
successful PLOGIs:    3
successful PRLIs:     3
successful INQUIRies: 3
successful LOGOs:     3
outstanding IUs:     0
probing state:        3
probing TOV:          0
probing count:        0
probing next:         0
pmap:                 0x00000000, 0x00000000, 0x00000000, 0x00000000
update map:           0x40000000, 0x00000000, 0x00000000, 0x00000000
npiv pmap:            0x00000000 0x00000000 0x00000000 0x00000000 \
                    0x00000000 0x00000000 0x00000000 0x00000000
npiv update map:     0x00000006 0x00000000 0x00000000 0x00000000 \
                    0x00000000 0x00000000 0x00000000 0x00000000
```

fcProbeShow

```
list of devices(may include old NPIV devices):  
0x500e8: SEAGATE ST318452FC      0005  
0x500e4: SEAGATE ST318452FC      0001  
0x500e2: SEAGATE ST318452FC      0005
```

SEE ALSO [portLoginShow](#), [portLogShow](#)

fcpRlsProbe

Initiates the Fibre Channel Protocol (FCP) Read Link Status (RLS) probing for F_Port and displays the RLS information.

Synopsis `fcpRlsProbe --start [slot/]port_list`

`fcpRlsProbe --show [slot/]port`

`fcpRlsProbe --help`

Description

Use this command to initiate the FCP RLS probing or to display the RLS information. This information describes the number of link failures, loss-of-signal, loss-of-sync, CRC errors, and other failure events detected on the specified port.

Use the `fcprlsprobe --start` command to read the error status block for F/FL-Ports. This command will send a RLS ELS probe to the device on the specified port and caches the RLS information.

Use the `fcprlsprobe --show` command to display the cached RLS information. For the command to send RLS probe to the device you must enable the following parameter: Disable Device Probing = 0 (Enables device probing. By default Device probing is enabled).

For this command to gather and display F_Port error statistics, apart from enabling the device probing, you must enable RLS Probing using the `fcprlsprobe --start` command or use the `configure` command, "Disable RLS Probing". By default, RLS probing is disabled.

Notes

The cached RLS information will be provided in `supportsave`. To have the updated RLS information for a specified port or port range in `supportsave`, execute `fcpRlsProbe --start` successfully before initiating the `supportsave`.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

| | |
|----------------------|--|
| <i>slot</i> | Specifies the slot number, if applicable, followed by a slash (/). |
| <i>port_list</i> | Specifies a port number. You can specify a port or range of ports. |
| <code>--start</code> | Sends the ELS RLS to the device and gets the RLS information. |
| <code>--show</code> | Displays the RLS information. |
| <code>--help</code> | Displays the command usage. |

Examples To start the RLS probing:

```
switch:admin> fcprlsprobe --start 1
RLS probing initiated...
Please use the fcprlsprobe --show port command to view the RLS info.
```

To view the RLS information:

```
switch:admin> fcprlsprobe --show 1
Rls Probing Statistics for port 1
=====
AL_PA 0x0
-----
RLS Probe Status      : SUCCESS
Last Updated         : Fri May 18 08:44:11 2012

-----
link fail loss sync loss sig prtc err bad word crc err
-----
      0      6      6      0      65555      0
```

SEE ALSO [fcpRlsShow](#)

fcpRlsShow

Displays the Fibre Channel Protocol (FCP) Read Link Status (RLS) information.

Synopsis `fcprlsshow [slot/]port`

Description Use this command to display the FCP RLS information for an F_Port or FL_Port. This information describes the number of loss-of-signal, loss-of-sync, CRC errors, and other failure events detected on the specified port.

For this command to gather and display F_Port error statistics,

- you must enable the configuration parameter: Disable Device Probing = 0 (Enables device probing. By default Device probing is enabled).
- you must enable RLS Probing in the following ways: Use the `fcprlsprobe --start` command which enables RLS Probing or use the `configure` command "Disable RLS Probing" to enable RLS probing. Disable RLS Probing = Off (Enables RLS Probing) by default RLS probing is disabled.

Use the `fcprlsprobe --show` command to view the RLS information.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to display, followed by a slash (/).

port

Specifies the port number to display, relative to its slot for bladed systems. Use `switchShow` for a listing of valid ports. This operand is required.

Examples To display the FCP RLS information:

```
switch:admin> fcprlsshow 9
RLS Probing Statistics for port 9
=====

AL_PA 0x1
=====
RLS Probe Status      : SUCCESS
Last Updated         : Mar 4 03:17:50 2012
-----
link fail loss sync  loss sig  prtc err  bad word  crc err
-----
1138      65047      1138      0          4996278  71

AL_PA 0x2
=====
RLS Probe Status      : FAIL (RLS Send timed out)
Last Updated         : Mar 4 03:17:50 2012
```

fcprIsShow

```
AL_PA 0x3
=====
RLS Probe Status      : SUCCESS
Last Updated         : Mar 4 03:17:52 2012
-----
link fail  loss sync  loss sig  prtc err  bad word  crc err
-----
147        1035      118       0         306278    72
```

SEE ALSO [portLoginShow](#), [portShow](#)

fcrBcastConfig

Displays or sets the broadcast frame forwarding option.

Synopsis

```
fcrbcastconfig --show
fcrbcastconfig --enable -f fabric_id
fcrbcastconfig --disable -f fabric_id
fcrbcastconfig --help
```

Description Use this command to enable or disable the broadcast frame option or to display the current configuration. If no operands are specified, this command displays the usage. By default, frame forward option is disabled. Use the **--show** option to display the current settings on the switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|---------------------|--|
| --show | Shows the current broadcast configuration as enabled or disabled. If broadcast frame forwarding is disabled for selected FIDs, only the enabled FIDs in the current configuration are displayed. |
| --enable | Enables the frame forwarding option for a specified fabric ID. |
| --disable | Disables the frame forwarding option for a specified fabric ID |
| -f fabric_id | Specifies the fabric ID to be disabled or enabled. Valid values are 1 to 128. This operand is required with the --enable and --disable options. |
| --help | Displays the command usage. |

Examples To display the current configuration:

```
fcr:admin> fcrbcastconfig --show
Broadcast configuration is disabled for all FID
```

To enable broadcast frame forwarding for FID 33, 28, and 2:

```
fcr:admin> fcrbcastconfig --enable -f 33
fcr:admin> fcrbcastconfig --enable -f 28
fcr:admin> fcrbcastconfig --enable -f 2
```

fcrBcastConfig

To display the new configuration:

```
fcr:admin> fcrbcastconfig --show  
Broadcast configuration is enabled for FID:  
2 33 128
```

To disable broadcast frame forwarding for FID 33:

```
fcr:admin> fcrbcastconfig --disable -f 33
```

To display the new configuration:

```
switch:admin> fcrbcastconfig --show  
Broadcast configuration is enabled for FID:  
2 128
```

SEE ALSO [bcastShow](#), [portRouteShow](#)

fcrConfigure

Displays or sets FC Router configuration parameters.

Synopsis

```
fcrconfigure --bbfid
fcrconfigure --enable -shortestifl
fcrconfigure --disable -shortestifl
fcrconfigure --add -alias Alias_name -fid FID
fcrconfigure --delete -alias [FID | -all]
fcrconfigure --resetphantomdomain [-force]
fcrconfigure --show [-alias]
fcrconfigure --help
```

Description Use this command to change the backbone fabric ID, or to enable or disable the shortest inter-fabric link (IFL) mode. If no operands are specified, this command displays the usage. Use the **--show** option to display the current settings on the switch.

Before you configure the backbone fabric ID using the **--bbfid** option, you must disable FC routing by using the **fosConfig** command and disable the switch using the **switchDisable** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--bbfid

Specifies the Backbone Fabric ID. This is an interactive option. A fabric ID uniquely identifies a fabric in FC Router configurations. The backbone fabric is the fabric attached to the U_Ports of the switch, for example, E_Ports or F_Ports. The backbone fabric ID must be unique across all fabrics connected to the FC Router. Use the **switchShow** command to display the current Backbone Fabric ID.

--enable -shortestifl

Enables the shortest IFL mode in FC Router.

When the shortest IFL mode is enabled, FC Router can choose a lowest-cost IFL path in the backbone fabric. This feature is useful when an FC Router has multiple connections to the source edge fabric and the backbone fabric has multiple FC Routers connected through FCIP links (VE_Ports) and FC links (E_Ports). Because a domain in an edge fabric can choose any equal cost path to reach the translate domain, frames can be transmitted through FCIP links even though FC links are present in the backbone fabric. When the shortest IFL mode is enabled, the FCIP path in the backbone fabric is avoided while reaching the destination edge fabric.

You should identify the FCIP links in the backbone fabric and then the direction of the FCIP path (across the backbone fabric) from the source to the destination edge fabric. Once the direction of the FCIP path is identified, set the cost of the FCIP link greater than or equal to 10000 by using the **linkCost** command. The link cost must be set for all FC Routers in the identified FCIP link path. Likewise, identify and set the cost for other FCIP paths of all destination edge fabrics.

After you identify the FCIP links between FC Routers for all fabrics and updated the cost of all FCIP links, enable the shortest IFL mode in all the FC Routers in the backbone fabric using this option. If the FCIP link is the only available path from the source to the destination edge fabric, then traffic will flow through that FCIP link only.

--disable -shortestifl

Disables the shortest IFL mode in FC Router.

--add -alias *Alias_name* -fid *FID*

Adds the alias name to the specific FID.

--delete -alias [*FID* | -all]

Deletes a specific FID alias association or all FID alias associations.

--resetphantomdomain

Resets the persistent front and translate domain IDs to the default starting value (160 for front domain ID and 200 for translate domain ID).

-force

Executes the command without confirmation. This operand is optional.

--show [-alias]

Shows the current Backbone Fabric ID with the associated alias name and the status of the shortest IFL mode as enabled or disabled.

--help

Displays the command usage.

Examples

To configure FC Router parameters:

```
fcr:admin> fcrconfigure --bbfid
Backbone fabric ID parameter set. <cr> to skip this parameter
Please make sure new Backbone Fabric ID does \
  not conflict with any configured EX-Port's Fabric ID
Backbone fabric ID: (1-128)[128]32
```

To enable shortest IFL mode:

```
fcr:admin> fcrconfigure --enable -shortestifl
Shortest IFL path is enabled.
```

To disable shortest IFL mode:

```
fcr:admin> fcrconfigure --disable -shortestifl
Shortest IFL path is disabled
```

To display the current configuration:

```
fcr:admin> fcrconfigure --show
Backbone fabric ID: 32
Shortest IFL feature is disabled
```

To assign a name to a FID:

```
fcr:admin> fcrconfigure --add -alias Red_fabric -fid 10
```

To display a name associated to a FID:

```
fcr:admin> fcrconfigure --show -alias
FIDAlias
=====
10Red_fabric
```

To delete the name of a FID:

```
fcr:admin> fcrconfigure --delete -alias 10
```

To reset all the phantom domain to the default range:

```
fcr:admin> fcrconfigure --resetphantomdomain
This operation will reset all the phantom domain to be default range
Do you want to continue (Y/N):y

Phantom Domain IDs were successfully reset to default
```

SEE ALSO [fosConfig](#), [portCfgEXPort](#), [switchDisable](#), [switchEnable](#), [switchShow](#), [linkCost](#)

fcrEdgeShow

Displays the FIDs of all configured EX_Ports.

Synopsis `fcredgeshow`

`fcredgeshow [-fid FabricID]`

`fcredgeshow --help`

Description

Use this command without operand to display information about all Fabric IDs (FIDs) that have been created on the chassis and are assigned to EX_Ports.

When a FID is specified, **fcredgeshow** displays information for all EX_Ports configured with the specified FID.

For each FID, the command output includes the following:

| | |
|----------------|--|
| FID | Fabric ID of the EX_Port. |
| EX-port | EX_Port number of the switch. |
| E_Port | Port number for the remote E_Port. |
| PWWN | Neighbor switch port WWN. |
| SWWN | Neighbor switch WWN. |
| Flags | encryption, compression and fec features |

The command output depends on the EX_Port configuration:

- If the EX_Port is online, the command displays the FID, the EX_Port to which it is assigned, the E_Port, the port WWN, the switch WWN and Port flags. Inter Chassis Link (ICL) EX_Ports do not support encryption and compression.
- If no EX_Ports are configured in the switch, the command displays "No EX-port Configured".
- If no EX_Ports are configured within the specified FID, the command displays the following message: "No EX-ports with FID *FabricID*."

Operands

This command has the following operands:

| | |
|-----------------------------|---|
| -fid <i>FabricID</i> | Specifies the FID for which to display the configured EX_Ports. |
| --help | Displays the command usage. |

Examples To display the EX_Ports configured in the switch:

```
switch:admin> fcredgeshow
FID  EX-port  E-port  Neighbor Switch (PWWN, SWWN ) Flags
-----
25   1/13      244    20:f4:00:05:1e:38:a4:cb 10:00:00:05:1e:38:a4:cb FEC \
      ENCRYPTION COMPRESSION
35   1/12      299    2e:2b:00:05:1e:40:44:02 10:00:00:05:1e:40:44:02 ENCRYPTION \
      COMPRESSION
11   5/13      273    2e:11:00:05:33:b3:39:00 10:00:00:05:33:b3:39:00 FEC
```

To display the EX_Ports configured with a specified FID:

```
switch:admin> fcredgeshow -fid 25
FID  EX-port  E-port  Neighbor Switch (PWWN, SWWN ) Flags
-----
25   11         244    20:f4:00:05:1e:38:a4:cb 10:00:00:05:1e:38:a4:cb FEC ENCRYPTION \
      COMPRESSION
```

To display a FID for which no EX_Ports are configured:

```
switch:admin> fcredgeshow -fid 29
No EX-ports with FID 29
```

SEE ALSO [fcrPhyDevShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [lsanZoneShow](#), [switchShow](#), [fcrFabricShow](#)

fcrFabricShow

Displays the FC Routers on a backbone fabric.

Synopsis `fcrfabricshow`

`fcrfabricshow --name`

`fcrfabricshow --help`

Description

Use this command to display information about FC Routers that exist in an FC Router backbone fabric. The existing syntax is maintained for IPv6 support. When IPv6 addresses are not configured, the output of **fcrFabricShow** displays the IPv4 format. Use the **--name** option to display the fabric name along with EX-port, FID, and switch name.

The message "No active FC Routers found" is displayed if no active FC Routers are present on the backbone fabric.

The following information is displayed for each FC Router found on the backbone fabric:

WWN

The world wide name of the FC Router.

Domain ID

The domain ID of the FC Router. This domain ID is relevant only on the backbone fabric.

Info

The Ethernet IP address and switch name of the FC Router. When IPv6 addresses are configured, only the static IP address displays for each FC Router found on the backbone fabric.

EX_Ports

A listing of active EX_Ports for the FC Router and information about these EX_Ports. This information includes:

EX_Port

The port number for the trunkmaster EX_Port. An asterisk (*) at the end of the line indicates that the EX_Port is a Remote Router Port.

FID

The fabric ID of the EX_Port.

Neighbor Switch Info (WWN, enet IP, name)

The WWN, Ethernet IP address, and switch name of the switch attached to the EX_Port.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operand.

--name

Displays the FC Routers on a backbone fabric with edge fabric names.

--help

Displays the command usage.

Examples To display the FC Routers in the backbone fabric:

```
switch:admin> fcrfabricshow

FC Router WWN: 10:00:00:05:1e:41:59:81, Dom ID: 2, Info:
10.33.36.8, "swd77"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)
-----
12      5 10.33.35.81 10:00:00:05:1e:34:01:d0 "B10_4"

FC Router WWN: 10:00:00:05:1e:41:1c:73, Dom ID: 4, Info:
10.33.36.12, "ttv12"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)
-----
9       2 10.33.35.80 10:00:00:05:1e:38:01:e7 "B10_3"
10      2 10.33.35.80 10:00:00:05:1e:38:01:e7 "B10_3"

FC Router WWN: 10:00:00:05:1e:39:51:67, Dom ID: 5, Info:
10.33.36.96, "Scimitar"
EX_Port FID Neighbor Switch Info (enet IP, WWN, name)
-----
151     2 10.33.35.80 10:00:00:05:1e:38:01:e7 "B10_3"*
```

To display the fabric name along with EX-port, FID, and switch name:

```
switch:admin> fcrfabricshow --name

FC Router WWN: 10:00:00:05:33:13:70:3e, Dom ID: 1,
Info: 10.17.33.129, "U34"
EX_Port FID Neighbor Switch Info (swname, fabricname)
-----
16      22 "SPIRIT-2" "FOSFCR"
17      22 "SPIRIT-2" "FOSFCR"
18      22 "SPIRIT-2" "FOSFCR"

FC Router WWN: 10:00:00:05:33:13:74:3e, Dom ID: 2,
Info: 10.17.33.128, "U35"
EX_Port FID Neighbor Switch Info (swname, fabricname)
-----
7       33 "STINGER-U33" "BODCFCR"
4       33 "STINGER-U33" "BODCFCR"
5       33 "STINGER-U33" "BODCFCR"
6       33 "STINGER-U33" "BODCFCR"
```

SEE ALSO [fcrPhyDevShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [lsanZoneShow](#), [switchShow](#)

fcrIcIPathBWMonitor

Sets or displays FC Router ICL bandwidth monitor parameters for EX_Ports.

Synopsis `fcrIcIpathbwmonitor --enable`

`fcrIcIpathbwmonitor --disable`

`fcrIcIpathbwmonitor --show`

`fcrIcIpathbwmonitor --help`

Description Use this command to enable or disable the Inter Chassis Link (ICL) EX_Ports bandwidth Monitor feature on an FC Router, or to display the current status of ICL bandwidth feature and the ICL slot bandwidth assigned to connected edge fabrics. If no operands are specified, this command displays the usage. By default, this feature is disabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--enable

Enables the ICL EX_Port bandwidth Monitor feature on an FC Router. In the enable mode, FC Router will periodically check for bandwidth imbalances from each ICL slot of FC Router to each ICL slot of neighbor switch connected through EX_Ports. A RASLog message will be generated when bandwidth imbalance is detected or resolved to other fabric.

--disable

Disables the ICL EX_Port bandwidth Monitor feature on an FC Router.

--show

Displays the current ICL EX_Port bandwidth Monitor feature status and bandwidth from each slot to connected edge fabric.

--help

Displays the command usage.

Examples To enable ICL EX_Port bandwidth Monitor:

```
fcr:admin> fcrIcIpathbwmonitor --enable
ICL bandwidth balance Monitor functionality is enabled
```

To disable ICL EX_Port bandwidth Monitor:

```
fcr:admin> fcrIcIpathbwmonitor --disable
ICL bandwidth balance Monitor functionality is disabled
```

To display the current configuration:

```
fcr:admin> fcrIclPathBWMonitor --show
ICL Path Bandwidth state :Enabled

FABRIC  SLOT-3 BW  SLOT-6 BW  STATE
=====
      48    128      128  BALANCED
      126    64      128  UNBALANCED
```

SEE ALSO [fosConfig](#), [switchDisable](#), [switchEnable](#), [switchShow](#)

fcrLsan

Configures and displays LSAN policies.

Synopsis

fcrIsan

fcrIsan --add -enforce tag | -speed tag

fcrIsan --remove -enforce tag | -speed tag

fcrIsan --show -enforce | -speed | -all

fcrIsan --help

Description

Use this command to add or remove LSAN tags, or to display existing tags in the configuration. LSAN tagging optimizes an FC Router's behavior based on a specified subset of LSANS. This feature improves scalability and performance related to LSAN zone size and the speed with which they are imported or exported.

This command supports two types of LSAN tags: enforced tags and speed tags.

- Enforced LSAN tags filter zones accepted by the FC Router from the edge fabric by matching the zones to the configured tags. Only matching zones are accepted into the local database for export and import. For example, if you configure an enforced LSAN tag "BRCD" on a router, only zones with names starting with "Isan_BRCD" are accepted. If multiple tags are configured, any matching zones are accepted. A maximum of eight LSAN enforce tags are configurable per FC Router switch.
- A speed tag is a flag to indicate to the FCR that the targets in the LSANs matching the tag need to be imported permanently when host and target are zoned together, even if the host is not present. This mechanism facilitates a speedy discovery process by reducing instances of failure related to timeouts. Once the devices that belong to the target edge fabric are defined as speed LSANS, the import or export can occur with a minimum amount of delay when hosts reboot or are added to the zone database.

The following restrictions apply when configuring LSAN tags:

- The switch must be disabled when you configure enforce tags. Speed tags can be configured while the switch is online.
- You must change the LSAN name in the edge fabric or the backbone fabric and propagate the LSAN to the FCR. Note that enforce tags are not supported in the backbone fabric.
- The speed tags must be set in all related FC Routers in order for import and export to proceed correctly. However, only LSANs on the target edge fabric must append the tag.

When executed without operands, **fcrIsan** displays the command usage.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

--add

Adds the specified tag to the LSAN tag configuration.

--remove

Removes the specified tag from the LSAN tag configuration.

| | |
|---------------------|---|
| --show | Displays the specified tag from the LSAN tag configuration. |
| --help | Displays the command usage. |
| -enforce tag | Accepts only the LSANs from the edge fabric that matches the specified tag string into the local FCR database. A valid tag is a string of a maximum of eight characters. The maximum configurable enforced tags is eight. |
| -speed tag | Allows the FCR to always import these target devices to the hosts specified in the LSANs that match the speed tag. Only one speed tag is allowed per FC router. |
| -all | When used with the --show option, this command displays all LSAN tags in the FCR LSAN tag database. |

Examples To add an LSAN enforcement tag named "brocade":

```
switch:admin> switchdisable

switch:admin> fcrLsan --add -enforce brocade
LSAN tag set successfully
```

To add a speed tag named "mcdt":

```
switch:admin> fcrLsan --add -speed mcdt
LSAN tag set successfully
```

To remove the LSAN enforcement tag "brocade":

```
switch:admin> fcrLsan --remove -enforce brocade
LSAN tag removed successfully
```

To remove the speed tag "mcdt":

```
switch:admin> fcrLsan --remove -speed mcdt
LSAN tag removed successfully
```

To display the information from the cache:

```
switch:admin> fcrLsan --show -enforce
Total LSAN tags : 2
ENFORCE : brocade
ENFORCE : cisco

switch:admin> fcrLsan --show -speed
Total LSAN tags : 1
SPEED: mcdt
```

fcrLsan

```
switch:admin> fcrLsan --show -all
Total LSAN tags : 3
ENFORCE : brocade
ENFORCE : cisco
SPEED: mcdt
```

SEE ALSO [fcrFabricShow](#), [lsanZoneShow](#), [fcrPhyDevShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [switchShow](#)

fcrLsanCount

Displays or sets the maximum LSAN count.

Synopsis `fcrlsancount [max_lsan_count]`

Description Use this command to set or display the maximum number of LSAN zones that can be configured on the edge fabric. By default, the maximum LSAN count is set to 3000, which is also the minimum. This command allows you to create up to 5000 LSAN zones in the backbone fabric, if needed to support additional devices. The maximum number of supported LSAN devices is 10,000.

When executed without operand, this command displays the current LSAN zone limit.

This command assumes that all FCRs in the same LSAN fabric matrix or backbone have the same maximum LSAN count defined in order to protect the FCRs from running into indefinite state. Asymmetric LSAN configurations due to different maximum LSAN counts may lead to different devices being imported on different FCRs.

Because the maximum number of LSANs is configured per switch, if there is a different maximum LSAN count on the switches throughout the meta-SAN, the device import or export will not be identical on the FCRs. You should therefore enter the same maximum LSAN count for all the FCR switches in the same backbone that support this feature. Verify the configured maximum limit against the LSANs configured using the **fcrResourceShow** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

max_lsan_count Specifies the maximum LSAN count.

Examples To display the current LSAN limit:

```
switch:admin> fcrlsancount
LSAN Zone Limit: 3000
```

To increase the LSAN zone limit:

```
switch:admin> fcrlsancount 5000
LSAN Zone Limit: 5000
```

SEE ALSO [fcrResourceShow](#)

fcrLsanMatrix

Creates, modifies, and displays the LSAN fabric matrix or the FCR matrix.

Synopsis

fcrlsanmatrix

fcrlsanmatrix --add -lsan *FID FID* | -fcr *wwn wwn*

fcrlsanmatrix --remove -lsan *FID FID* | -fcr *wwn wwn*

fcrlsanmatrix --apply -lsan | -fcr | -all

fcrlsanmatrix --cancel -lsan | -fcr | -all

fcrlsanmatrix --display -lsan | -fcr

fcrlsanmatrix --fabricview -lsan | -fcr

fcrlsanmatrix --verify -lsan | -fcr

fcrlsanmatrix --quickmode -lsan | -fcr

fcrlsanmatrix --help

Description

Use this command to create, modify, remove, and manage Logical SAN (LSAN) Zone bindings between edge fabrics or between FC routers. LSAN Zone bindings specify pairs of edge fabrics or FCRs that can access each other and share LSAN Zone and device database information.

This command follows a transaction model. Modifications to the LSAN fabric matrix or to the FC router matrix are saved nonpersistently in a cache until you save the changes persistently with the **--apply** option.

When used with the **-lsan** option, this command manages the LSAN matrix information. An LSAN fabric pair binds two edge fabrics specified by their Fabric IDs. Every paired edge fabric implies two-way communications. The paired edge fabrics have access only to the edge fabrics associated with them by this command. The edge fabrics that are not specified in the LSAN fabric matrix have access to the remaining unspecified edge fabrics. Using this information, the FCR switch maintains the remote LSAN Zone and the device state database only if it is associated with its local edge fabrics.

For example, if the edge fabrics with FIDs 1, 2, 3, 4, and 5 are online, all edge fabrics have two-way communication. This is the default behavior. If you pair the edge fabrics 1 and 2 with the **--add -lsan** command, the default access between the edge fabrics is changed as follows:

- Fabric 1 can access only fabric 2.
- Fabric 2 can access only fabric 1.
- The fabrics 3, 4, and 5 can access each other, but cannot access Fabric 1 or 2.

The LSAN matrix information is automatically distributed to all switches in the fabric. The FIDs entered are not required to be online when you set up the LSAN fabric matrix.

When used with the **-fcr** option, this command manages the FC router matrix. This database consists of FC router pairs that can talk to each other. All edge fabrics connected to a defined pair of FCRs are allowed to import devices to each other.

Once a fabric is removed from an FCR, the communication with other fabrics of the two FCRs is also removed. It is advisable to update the matrix to reflect the change. If FCR Binding is enabled in the edge fabrics, the edge fabrics can still communicate with the backbone fabric. The LSAN matrix provides a higher level of granularity than the FCR matrix. Therefore, communication between two fabrics could pass the FCR matrix, but fail due to restriction of the FID matrix.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands If no operands are specified, this command displays the persistent LSAN Zone matrix information. The following operands are optional:

--add -lsan *FID FID* | -fcr *wwn wwn*

Adds the pair of edge fabrics or FCR members that can access each other to the LSAN matrix cache. Edge fabrics are specified by their Fabric IDs. FCRs are specified by their world wide names (WWNs). If the FCR is online, the domain ID of the switch can be used in place of the WWN. If you specify **--add** with zero (0) value for *FID* or 00:00:00:00:00:00:00:00 for *wwn*, the command returns the cache to default mode.

--remove -lsan *FID FID* | -fcr *wwn wwn*

Removes the specified pair of FIDs or FCRs from the cache. When you remove a fabric or FCR pair from the LSAN matrix, the edge fabrics assume the default behavior.

--apply -lsan | -fcr | -all

Applies the information from the cache to persistent memory. This operation succeeds only if there is no effect on the existing import/export devices. Otherwise, this command generates an error message. See the Diagnostics section for more information.

--cancel -lsan | -fcr

Cancels changes made to the cache that were not applied. This command effectively reverts to the persistently stored information.

--display -lsan | -fcr

Displays the information saved in the cache.

--fabricview -lsan | -fcr

Displays all static and default/dynamic fabric bindings in the backbone.

--verify -lsan | -fcr

Verifies if the LSAN Zone information previously entered and stored in the cache can be successfully applied. The data is considered acceptable if the apply operation does not cause any traffic disruption

--quickmode

Derives the LSAN Zone matrix from the imported or exported devices.

Diagnostics Error message (1)

LSAN Matrix in the cache conflicts with existing import/export devices and may disrupt traffic. Please refer to the man page for the corrective action.

Corrective actions:

- Any new router added to the backbone fabric automatically triggers a matrix merge. If a router does not support the matrix merge feature, the router cannot join the backbone fabric. Make sure that all legacy FCR switches in the backbone support the matrix merge feature, otherwise the feature is not supported.
- Use **fcrLsanmatrix --fabricview -lsan | -fcr** to confirm that all the switches in the backbone have the same LSAN and FCR Binding matrix. If not, there are two solutions. The first solution is to modify one FCR or both to make them the same and then activate the FCRs. The second solution is to zero out the database of one FCR to signal that this FCR accepts the database from the other FCR once the change is activated.

To zero out the database, execute the following commands:

- **fcrLsanmatrix --add -lsan 0 0**
- **fcrLsanmatrix --add -fcr 00:00:00:00:00:00:00:00 00:00:00:00:00:00:00:00**
- **fcrLsanmatrix --apply -all**
- In a dual backbone configuration, execute **fcrLsanmatrix --fabricview** on the FCR switches to confirm that the shared edge fabric FIDs have the same access in both backbones.
- Execute **fcrLsanmatrix --display -lsan | -fcr** and **fcrproxydevshow -a**. Check that the LSAN Binding matrix in the cache is not in conflict with the existing import/export devices that are displayed on the FCR switch. If there is a conflict, do one of the following:
 - Update the LSAN/FCR Binding matrix in the cache to allow access for the FIDs that have imported devices.
 - Remove the conflicting import/export devices by updating the LSAN zone in the edge fabrics.
 - Disable the conflicting devices.

Error message (2)

There may be other FCR switches in the backbone that do not support the LSAN Binding feature or do not have the same fcrLsanmatrix settings.

Please refer to the man page for the corrective action.

Corrective actions:

- Verify that all FCR switches in the backbone support the LSAN Binding features.
- Execute **fcrLsanmatrix --fabricview** to confirm that all the switches in the backbone have the same LSAN Binding matrix. If not, clear the LSAN Binding feature on all the switches and reapply the same LSAN Binding matrix on all the FCR switches in the backbone.
- In a dual backbone configuration, use **fcrLsanmatrix --fabricview** on the FCR switches to confirm that the shared edge fabric FIDs have the same access in both backbones.

Refer to the *Fabric OS Message Reference manual* for further diagnostic information.

Examples

To add the LSAN Zone Matrix data (For the following example, assume that the backbone has the following online edge fabrics (FIDs): 1, 2, 4, 5, 7, 8, and 10. Currently, FIDs 14 and 19 are not available.):

```
switch:admin> fcrLsanmatrix --add -lsan 4 5
switch:admin> fcrLsanmatrix --add -lsan 4 7
switch:admin> fcrLsanmatrix --add -lsan 10 14
switch:admin> fcrLsanmatrix --add -lsan 10 19
```

To remove an entry from the LSAN matrix:

```
switch:admin> fcrLsanmatrix --remove -lsan 10 14
```

To display the information from the cache:

```
switch:admin> fcrlsanmatrix --display -lsan
```

| Fabric ID 1 | Fabric ID 2 |
|-------------|-------------|
| 4 | 5 |
| 4 | 7 |
| 10 | 19 |

To apply the changes persistently:

```
switch:admin> fcrlsanmatrix --apply -lsan
```

To view the persistent changes:

```
switch:admin> fcrlsanmatrix -lsan
LSAN MATRIX is activated
Fabric ID 1          Fabric ID 2
-----
4                   5
4                   7
10                  19
```

To view the LSAN Zone static and default/dynamic binding in the backbone where online fabrics are: 1, 2, 4, 5, 7, 8, 10:

```
switch:admin> fcrlsanmatrix --fabricview -lsan
LSAN MATRIX is activated

Fabric ID 1          Fabric ID 2
-----
4                   5
4                   7
10                  19
Default LSAN Matrix:
1 2 8
```

To display all proxy devices for all FC Routers in the same backbone fabric whether or not they are relevant to this FC Router:

```
switch:admin> fcrproxydevshow -a
```

| Proxy Created in Fabric | WWN | Proxy PID | Device Exists in Fabric | Physical PID | State |
|-------------------------|-------------------------|-----------|-------------------------|--------------|----------|
| 52 | 10:00:00:06:2b:0e:4d:e5 | 01f001 | 78 | 4e0000 | Imported |
| 52 | 10:32:16:90:28:dd:d0:03 | 0bf001 | 82 | 2a0900 | Imported |
| 52 | 10:32:16:91:24:dd:d0:07 | 0bf002 | 82 | 520c00 | Imported |
| 52 | 10:32:16:91:25:dd:d0:06 | 01f002 | 78 | 4e3000 | Imported |
| 78 | 10:00:00:06:2b:0d:29:31 | 09f002 | 52 | 482200 | Imported |
| 78 | 10:32:16:90:29:dd:d0:07 | 08f002 | 82 | 2a0a00 | Imported |
| 78 | 10:32:16:91:24:dd:d0:05 | 09f001 | 52 | 48a100 | Imported |
| 78 | 10:32:16:91:25:dd:d0:03 | 08f001 | 82 | 520f00 | Imported |

```

      82 10:00:00:06:2b:0d:29:30 01f002 78 4e1400 Imported
      82 10:00:00:06:2b:0d:2f:ed 03f002 52 480200 Imported
      82 10:00:00:06:2b:0d:33:4d 01f001 78 4e1800 Imported
      82 10:00:00:06:2b:0e:4d:c9 03f001 52 482000 Imported
Total devices displayed: 12

```

To display the information from the cache:

```
switch:admin> fcrLsanmatrix --display -lsan
```

```

Fabric ID 1          Fabric ID 2
-----
52              78
52              82
78              82

```

To apply the changes persistently:

```
switch:admin> fcrLsanmatrix --apply -lsan
```

To view all the static and the default/dynamic fabric binding in the backbone:

```
switch:admin> fcrLsanmatrix --fabricview -lsan
LSAN MATRIX is activated
```

```

Fabric ID 1          Fabric ID 2
-----
52              78
52              82
78              82

```

```

Default LSAN Matrix:
57 91

```

To add FCR Bindings to the FCR matrix:

```

switch:admin> fcrLsanmatrix --add -fcr 10:00:00:60:69:e2:09:fa \
10:00:00:60:69:e2:09:fb

switch:admin> fcrLsanmatrix --add -fcr 10:00:00:60:69:e2:09:fb \
10:00:00:60:69:e2:09:fc

```

To remove an entry from the FCR matrix:

```
switch:admin> fcrLsanmatrix --remove -fcr 10:00:00:60:69:e2:09:fb \
10:00:00:60:69:e2:09:fc
```

To display the information from the cache:

```
switch:admin> fcrLsanmatrix --display --fcr
```

```

CACHE FCR PAIRS
=====
FCR              FCR
-----
10:00:00:60:69:e2:09:fa (2) 10:00:00:60:69:e2:09:fb (unknown)

```

To apply the changes persistently:

```
switch:admin> fcrLsanmatrix --apply -fcr
```

To view the persistent changes:

```
switch:admin> fcrLsanmatrix --fabricview -fcr
```

```
          SAVED FCR PAIRS
=====
FCR                               FCR
-----
10:00:00:60:69:e2:09:fa (2)       10:00:00:60:69:e2:09:fb (unknown)
```

SEE ALSO [fcrFabricShow](#), [lsanZoneShow](#), [fcrPhyDevShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [switchShow](#)

fcrPhyDevShow

Displays the FC Router physical device information.

Synopsis `fcrphydevshow [-a][-f FID][-w wwn][-c][-d] [-h]`

Description Use this command to display the physical (real) devices that are configured to be exported to other fabrics. A device is considered to be configured to be exported to another fabric if it is a member of an LSAN zone. The device is displayed only if it is discovered in the EX_Port-attached fabric and backbone fabric's name server (for instance, the device is online).

Physical device information is available only for physical devices that exist in fabrics attached to EX_Ports of FC Routers on the same backbone fabric as the current FC Router.

The default output displays only physical device information relevant to the current FC Router. Relevant physical devices include physical devices that are configured to be exported from fabrics attached to the current FC Router's EX_Ports.

The physical devices are listed by fabric.

The **-f** and **-w** operands allow searching for physical devices based on fabric ID or port world wide name.

"No device found" is displayed if there is no physical device information available at the current FC Router.

Each line of the output displays:

Device Exists in Fabric

The fabric in which the physical device exists.

WWN

The world wide name of the device port.

Physical PID

The port ID of the physical device. This port ID is only relevant on the fabric specified by the "Device Exists in Fabric" column.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-a

Displays all physical devices for all FC Routers in the same backbone fabric whether or not they are relevant to the current FC Router.

-a -f *FID*

Displays the physical devices in the specified fabric for all FC Routers in the same backbone fabric whether or not they are relevant to the current FC Router.

-w *wwn*

Displays the physical devices with the specified port WWN.

-c

Clears login-related counters.

-d Displays the following login-related counters. Counters are cleared upon reboot or failover.

login try The number of times the device attempted to log in.

local failure the number of times the device login failed because of missing LSAN zones within the device fabric.

remote failure the number of times the device login failed due to missing LSAN zones within the remote fabric.

-h Displays command usage.

Examples To display the physical devices relevant to an FC Router:

```
fcr:admin> fcrphydevshow
      Device           WWN           Physical
      Exists
      in Fabric
-----
      2    10:00:00:00:c9:2b:6a:68  c70000
      3    50:05:07:65:05:84:09:0e  0100ef
      3    50:05:07:65:05:84:0b:83  0100e8
Total devices displayed: 3
```

SEE ALSO [fcrFabricShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [lsanZoneShow](#), [switchShow](#)

fcrProxyConfig

Displays or configures proxy devices presented by an FC Router.

Synopsis `fcrproxyconfig`
`[-s importedFID devWWN slot]`
`[-r importedFID devWWN]`

Description Use this command to display or set the persistent configuration of proxy devices presented by the local FC Router.

When used without operand, this command displays the persistent proxy device configuration; otherwise, it sets the specified attributes to its new value.

The proxy device must be inactive prior to setting or clearing persistent attributes. Disabling EX_Ports (using the **portDisable** command) attached to the relevant edge fabric, removing the device from the appropriate LSAN zones, or disabling the physical device are valid methods of ensuring a proxy device is inactive.

Persistent proxy device configuration attributes apply to the local FC Router. Multiple FC Routers attached to the same edge fabric coordinate to present the same proxy devices. As a result, persistent proxy device configurations must be consistent across all FC Routers attached to the same edge fabric or unpredictable results may occur. If the proxy device configuration is not altered, no action is required. If the configuration is altered, then care must be taken to ensure consistency across all FC Routers attached to the same edge fabric.

When used without operands, this command displays the following information:

importedFID

The imported fabric ID of the proxy device.

devWWN

The port world wide name of the device.

Slot

The slot used for the device WWN. The device WWN-to-slot association is persistently stored. The slot format is *XXYYH*, where *XX* specifies the translate domain area_ID (valid values include FOH through FFH) and *YY* specifies the Port ID value or the low 8-bits of the proxy device address (valid values include 01H through 7FH). The address of the proxy device is derived from the PID format (for example, native, core, or extended edge) and the proxy device slot.

If no proxy device WWN is stored in any slot for all edge fabrics, the following message is displayed: "All slots empty."

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-s *importedFID devWWN slot*

Adds the specified *devWWN* (format: xx:xx:xx:xx:xx:xx:xx) to the specified slot (format XXYH, where XX is the translate domain area_ID [FOH through FFH] and YY is the port_ID [01H through 7FH]) for the edge fabric specified (1 through 128). The following messages may be displayed:

- "WWN does not exist in any proxy device slot." The WWN does not exist in any slot for the specified edge fabric.
- "Too many proxy slots configured. Remove some unused proxy device WWNs from their slots using the -r operand and try again." All slots are used for the specified edge fabric.
- "The specified slot already contains a WWN, overwrite? [y]" The specified slot already contains an entry. You are then prompted for confirmation.

-r *importedFID devWWN*

Removes the specified *devWWN* (format: XX:XX:XX:XX:XX:XX:XX) from its slot for the edge fabric specified by *importedFID* (1 through 128). If the WWN does not exist in any slot for the specified edge fabric, the following message is displayed: "WWN does not exist in any proxy device slot."

Examples To display the persistent proxy device configuration:

```
switch:admin> fcrproxyconfig
Imported FID      Device WWN          Slot
002              50:05:07:65:05:84:d7  f001
002              50:05:07:65:05:84:0a:7b f002
002              22:00:00:20:37:c3:11:71 f001
002              22:00:00:20:37:c3:1a:8a f002
003              10:00:00:00:c9:2b:6a:2c f001
```

To persistently configure device WWN 00:11:22:33:44:55:66:77 to use slot f101h in fabric 5:

```
switch:admin> fcrproxyconfig -s 5 \
00:11:22:33:44:55:66:77 f101
```

To remove device WWN 00:11:22:33:44:55:66:77 from its persistent slot in fabric 5:

```
switch:admin> fcrproxyconfig -r 5 \
00:11:22:33:44:55:66:77
WWN deleted from proxy device slot
```

SEE ALSO [fcrPhyDevShow](#), [fcrProxyDevShow](#), [fcrXlateConfig](#), [lsanZoneShow](#), [switchShow](#)

fcrProxyDevShow

Displays FC Router proxy device information.

Synopsis `fcrproxydevshow [-a][-f fabricid][-w wwn]`

Description Use this command to display the proxy devices presented by FC Router EX_Ports and information about the proxy devices. A proxy device is a virtual device presented in to a fabric by an FC Router. A proxy device represents a real device on another fabric. When a proxy device is created in a fabric, the real Fibre Channel device is considered to be imported in to this fabric. The presence of a proxy device is required for inter-fabric device communication. The proxy device appears to the fabric as a real Fibre Channel device. It has a name server entry and is assigned a valid port ID.

Proxy device information is available only for proxy devices that are presented by FC Routers on the same backbone fabric as this FC Router.

The default output displays only proxy device information relevant to this FC Router. Relevant proxy devices include proxy devices created by this FC Router (devices imported by this FC Router).

The proxy devices are listed by fabric. Search parameters **-f** and **-w** allow searching for proxy devices based on fabric ID or port WWN.

"No proxy device found" is displayed if there is no proxy device information available on this FC Router.

This command displays the following information:

Proxy Created in Fabric

The fabric in which the proxy device has been created.

WWN

The WWN of the device port.

Proxy PID

The port ID of the proxy device. The port ID is only relevant on the fabric specified by the "Proxy Created in Fabric" column.

Device Exists in Fabric

The fabric in which the physical device represented by this proxy device exists.

Physical PID

The port ID of the physical device. The port ID is relevant only on the fabric specified by the "Device Exists in Fabric" column.

State

State includes:

Imported

Proxy device has been imported into the fabric.

Initializing

The proxy device is being initialized and will soon be imported into the fabric.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

- a** Display all proxy devices for all FC Routers in the same backbone fabric whether or not they are relevant to this FC Router.
- a -f fabricid** Display the proxy devices in the specified fabric for all FC Routers in the same backbone fabric whether or not they are relevant to this FC Router.
- f fabricid** Display the proxy devices in the specified fabric that are relevant to this FC Router.
- w wwn** Displays proxy devices with the specified port WWN.

Examples To display the proxy devices relevant to this FC Router:

```
switch:admin> fcrproxydevshow
Proxy          WWN          Proxy  Device  Physical State
Created                               PID    exists  PID
in Fabric                               in Fabric
-----
2   50:05:07:65:05:84:09:0e 01f001    3   0100ef Imported
2   50:05:07:65:05:84:0b:83 01f000    3   0100e8 Imported
3   10:00:00:00:c9:2b:6a:68 02f000    2   c70000 Imported
Total devices displayed: 3
```

SEE ALSO [fcrFabricShow](#), [fcrRouteShow](#), [lsanZoneShow](#), [switchShow](#)

fcrResourceShow

Displays FC Router physical resource usage.

Synopsis `fcrresourceshow`

Description Use this command to display the FC Router-available resources. The maximum number allowed versus the currently used is displayed for various resources. The command output includes:

LSAN Zones

The maximum versus the currently used LSAN zones.

LSAN Devices

The maximum versus the currently used LSAN device database entries. Each proxy or physical device constitutes an entry.

Proxy Device Slots

The maximum versus the currently used proxy device slots. A proxy device is presented to an edge fabric as being connected to a translate domain slot. A slot is the port number and AL_PA combination. The slot-to-device WWN association is persistently stored.

Phantom Node WWNs

The maximum versus the currently allocated phantom switch node WWNs. The phantom switch requires node WWNs for fabric-shortest-path-first (FSPF) and manageability purposes. Phantom node names are allocated from the pool sequentially and are not reused until the pool is exhausted and rolls over. The last allocated phantom node WWN is persistently stored. If the switch is disabled, the phantom node WWNs are not returned to the pool because the phantom switch could still be accessible through other switches. Across a switch reboot, the allocation starts from the next usable WWN from the pool and not from the beginning.

Phantom Port WWNs

The maximum versus the currently used phantom domain port WWNs. Phantom domain ports require port WWNs for manageability purposes. Phantom domain ports include ports connecting front and translate domains (virtual ISLs), translate domain ports for proxy devices, and EX_Ports. Phantom port names are allocated from the pool sequentially and are not resumed until the pool is exhausted and rolls over. The last allocated phantom port WWN is persistently stored. If the switch is disabled, phantom port WWNs are not returned to the pool because the phantom switch might still be accessible through other switches. Across the switch reboot, the allocation starts from the next usable WWN base from the pool and not from the beginning.

Port Limits

Displays resources for each physical port (EX_Port), which include the following:

Max Proxy Devices

The maximum versus the currently used proxy device.

Max NR_Ports

The maximum versus the currently used NR_Port entries. Destination NR_Port entries are stored at every physical port for routing decision purposes.

Notes Only configured EX/VEX_Ports are displayed

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the resource usage for the local FC Router:

```
switch:admin> fcrresourceshow

      Daemon Limits:
              Max Allowed  Currently Used
      -----
LSAN Zones:           3000                22
LSAN Devices:        10000             1208
Proxy Device Slots: 10000                2

              WWN Pool Size  Allocated
      -----
Phantom Node WWN:    8192             3790
Phantom Port WWN:   32768             6446

Port Limits:
Max proxy devices:  2000
Max NR_Ports:      1000

Currently Used(column 1: proxy, column 2: NR_Ports):
48 |      0      0
49 |      0      0
50 |      0      0
52 |      0      0
53 |      0      0
54 |      0      0
60 |      0      0
63 |      1      4
176 |     1      4
177 |     1      4
183 |     1      4
190 |      0      0
```

SEE ALSO [fcrFabricShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [lsanZoneShow](#), [switchShow](#)

fcrRouterPortCost

Displays or sets an FC Router port cost.

Synopsis `fcrrouterportcost [[slot/]port] [cost]`

Description Use this command to set or display the cost of the FC Router ports. You can set the cost of the link to one of two fixed values: 1000 or 10000. The option 0 sets the cost of the link to the default value based on link type (EX/VEX). The router module chooses the router port path based on the minimum cost per fabric ID (FID) connection. If multiple paths exist with the same minimum cost, the load is shared over these paths.

Every inter-fabric link (IFL) has a default cost. For an EX_Port IFL, the default cost is 1000. For a VEX_Port, the default cost is 10000. If the cost is set to 0, the link cost defaults to 1000 for an EX_Port and to 10000 for a VEX_Port.

when used without operands, this command displays the current link costs for all ports on the switch.

Notes Before setting the cost, ensure that admin is enabled for the EX_Port/VEX_Port with **portCfgEXPort** or **portCfgVEXPort**. The cost can be set only on a disabled port.

The bandwidth of an inter-fabric link (IFL) is unrelated to its default cost. In other words, 1 Gbps, 2 Gbps, 4 Gbps, and 8 Gbps EX_Port IFLs have the same cost value of 1000 as their FC Router port.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port whose cost is to be displayed or changed, followed by a slash (/).

port

Specifies the number of the port whose cost is to be displayed or changed. This value is relative to the slot for bladed systems. Use **switchShow** for a list of valid ports. If a port is not specified, this command displays the costs of all ports.

cost

Specifies the new cost of the link connected to the specified port. This operand is optional; if omitted, this command displays the cost of the specified port. The cost of the link can be changed only when the specified port is disabled. Valid values for cost are 0, 1000 or 10000.

Examples To display the cost of all EX_Ports:

```
switch:admin> fcrrouterportcost
Port          Cost
-----
7/3           1000
7/4           1000
7/9           1000
7/10          1000
7/16          10000
10/0          10000
```

To display the cost on an EX_Port:

```
switch:admin> fcrrouterportcost 7/10 0

switch:admin> fcrrouterportcost 7/10
Port          Cost
-----
7/10          1000
```

To set the cost of an EX_Port and display the result:

```
switch:admin> fcrrouterportcost 7/10 10000
switch:admin> fcrrouterportcost 7/10
Port          Cost
-----
7/10          10000
```

To set the default cost on the EX_Port:

```
switch:admin> fcrrouterportcost 7/10 0
switch:admin> fcrrouterportcost 7/10
Port          Cost
-----
7/10          1000
```

SEE ALSO [switchShow](#), [fcrRouteShow](#), [portCfgEXPort](#)

fcrRouteShow

Displays FC Router route information.

Synopsis `fcrrouteshow`

Description Use this command to display routes through the FC Router backbone fabric to accessible destination fabrics. An FC Router backbone fabric is the fabric that contains the E_Ports of this platform and routes inter-fabric traffic between imported fabrics, creating a meta-SAN.

There are FC Router ports that reside on the backbone fabric. These ports are known as NR_Ports. NR_Ports send and receive inter-fabric traffic. For the AP7420, there is a one-to-one relationship between an NR_Port on a backbone fabric and an EX_Port. NR_Port technology enables EX_Ports to exchange traffic across an intermediate fabric. NR_Ports are addressable entities on the backbone fabric and have port IDs relevant to the backbone fabric.

Because cascaded backbone/intermediate fabrics are currently not supported, an NR_Port provides a path to a single fabric with a single FC Router protocol cost. Multiple NR_Ports can provide paths to the same destination fabric.

"No routes found" is displayed if there is no route information available at this FC Router. There is no route information available if no EX_Ports are configured at this FC Router.

The output includes:

Destination Fabric ID

The destination fabric.

NR_Port PID

The port ID of the NR_Port. The port ID is relevant only on the backbone fabric. This NR_Port has a route to the destination fabric identified by the "Destination Fabric ID" column.

FCRP Cost

The FC Router protocol cost (for routing decisions) for this NR_Port. The FCRP cost is the same (1000) for all NR_Ports.

WWN of the Principal Switch in the Dest. Fabric

The world wide name of the principal switch in the destination fabric.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the route information:

```
switch:admin> fcrrouteshow
Destination NR_Port   FCRP Cost  WNN of Principal
Fabric Id   PID           Switch in the Dest. Fabric
-----
      4      640000      1000  10:00:00:60:69:c0:05:d1
      4      640100      1000  10:00:00:60:69:c0:05:d1
      5      640200      1000  10:00:00:60:69:c0:20:ed
      5      640300      1000  10:00:00:60:69:c0:20:ed
```

SEE ALSO [fcrFabricShow](#), [fcrPhyDevShow](#), [fcrProxyDevShow](#), [lsanZoneShow](#), [switchShow](#)

fcrXlateConfig

Configures a translate (xlate) domain's domain ID and state of persistence for both the EX_Port-attached fabric and the backbone fabric.

Synopsis

fcrxlateconfig

fcrxlateconfig *importedFID exportedFID preferredDomainID*

fcrxlateconfig --remove | **-r** *importedFID exportedFID*

fcrxlateconfig --enable **persistxd**

fcrxlateconfig --disable **persistxd**

fcrxlateconfig --show **stalex** *importedFID*

fcrxlateconfig --delete **stalex** *importedFID staleXlateDomainID*

fcrxlateconfig --help

Description

Use this command to display a translate (xlate) domain ID or change the preferred domain ID and its state of persistence.

A translate domain is a phantom domain created by an FC Router. FC Routers emulate proxy devices representing real devices in remote fabrics. These proxy devices are emulated to be connected to translate domains. Translate domains are presented to a fabric as residing topologically behind front phantom domains (domains created by an EX_Port). In the case of backbone fabrics, translate domains are topologically behind an E_Port. In every EX_Port-attached edge fabric and backbone fabric, there can be a translate domain for every FC Router-accessible remote fabric.

During a fabric build, the translate domain requests a domain ID from the principal switch in the EX_Port-attached edge fabric. The domain ID requested is the preferred domain ID. You can set the preferred domain ID when the translate domain is not active and is persistently saved. The principal switch attempts to provide the translate domain with the requested domain ID, but it may not provide it if there are domain ID conflicts with other domains in the fabric. If the requested domain ID (such as the preferred domain ID) is unavailable, the domain ID assignment is completely at the discretion of the principal switch. The assignment domain ID is persistently stored and is used as the preferred domain ID in the future.

By default, FCR creates the translate domain for a remote fabric if a valid persistent translate domain ID is configured in the local fabric, even if no devices are imported or exported across the edge fabrics. Disabling the **persistxd** parameter prevents the xlate domain from being created. Enabling the **persistxd** parameter from a disabled state re-enables the FCR default behavior.

If the remote edge fabric becomes unreachable, the translate domains created in other edge fabrics for this remote edge fabric become stale. Use the **--show stalex** option to identify translate domains that have become stale, and delete them in a nondisruptive manner with the **--delete stalex** option.

When executed without operands, **fcrxlateconfig** displays for each translate domain the imported FID, the exported FID, the domain ID and the xlate WWN.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Information displayed is not related to the entire backbone. The FC Router displays only connections to an edge fabric for which there are translate domain IDs. Any changes you intend to make using this command should be issued on the switches to which the edge fabrics are directly attached. In a Virtual Fabric environment, this is the base switch.

Operands This command has the following operands:

fcrxlateconfig

Sets the preferred domain ID (1-239) to *preferredDomainID* for the translate phantom domain and saves the configuration persistently. The translate domain must be inactive to set the preferred domain ID. The following operands are required:

importedFID

Specifies the fabric ID (1 through 128) of the fabric that contains the translate domain.

exportedFID

Specifies the fabric ID (1 through 128) of the remote fabric represented by this translate domain.

preferredDomainID

Specifies the preferred domain ID (1 through 239) of the translate phantom domain.

--remove | -r

Removes the preferred domain ID of the translate phantom domain. The translate domain must be inactive to remove the preferred domain ID. The following operands are required:

importedFID

Specifies the fabric ID (1 through 128) of the fabric that contains the translate domain.

exportedFID

Specifies the fabric ID (1 through 128) of the remote fabric represented by this translate domain.

preferredDomainID

Specifies the preferred domain ID (1 through 239) of the translate phantom.

--enable persistxd

Enables translate domain persistence. When **persistxd** is enabled, the translate domain is created based on the persistent translate domain ID configuration. If a valid persistent translate domain ID is configured for a given *importedFID* and *exportedFID* pair, a translate domain for the *exportedFID* is created, even if no devices need to be imported or exported across the edge fabrics represented by *importedFID* and *exportedFID*. By default, **persistxd** is enabled.

--disable persistxd

Disables translate domain persistence. When **persistxd** is disabled, the translate domain is not created, even if a valid persistent translate domain ID is configured for the *importedFID* and *exportedFID* pair, so long as no devices are imported or exported across the edge fabrics represented by *importedFID* and *exportedFID*. Once devices need to be imported or exported across the edge fabrics, the translate domain is created.

--show stalexd [importedFID]

Displays stale translate domains associated with the specified Fabric ID (1-128). A translate domain becomes stale when the remote edge fabric for which this translate domain was created in the specified edge fabric becomes unreachable. When issued without specifying an imported FID, this command lists all stale translate domains in all edge fabrics connected to the FCR.

--delete stalexd importedFID staleXlateDomainID

Deletes the specified stale translate domain from the edge fabric specified by its fabric ID. This command must be executed in the FCR that owns the stale translate domain.

Examples

To display the translate domain configuration and the state of the **persistxd** parameter:

```
switch:admin> fcrxlateconfig

ImportedFid ExportedFid Domain OwnerDid XlateWWN
001 002 004 000001 N/A
001 005 003 N/A N/A

Persist XD state: Enabled
```

To set the preferred domain ID of the translate domain created in fabric 2, which represents the remote fabric 3, to a value of 8:

```
switch:admin> fcrxlateconfig 2 3 8
xlate domain already configured, overwrite?(n) y
```

To clear the preferred domain ID of the translate domain created in fabric 2, which represents remote fabric 3:

```
switch:admin> fcrxlateconfig -r 2 3
xlate domain deleted
```

To enable translate domain persistence:

```
fcr:admin> fcrxlateconfig --enable persistxd
Persist XD is enabled
```

To disable translate domain persistence:

```
fcr:admin> fcrxlateconfig --disable persistxd
Persist XD is disabled
```

To identify and remove stale translate domains in a single backbone multiple FCR configuration:

```
fcr:admin> fcrxlateconfig --show stalexd
Imported FID      Stale XD      Owner Domain
-----
012              002           007 ( this FCR )
013              002           001 ( other FCR )
```

To remove the stale translate domain (only the translate domain owned by the current owner can be removed; note that trailing zeros must be removed from the FID and Xlate domain ID.>):

```
fcr:admin> fcrxlateconfig --delete stalexd 12 2
Xlate domain 2 is deleted
```

SEE ALSO [portCfgEXPort](#), [portCfgVEXPort](#), [portDisable](#), [portEnable](#), [portShow](#)

fddCfg

Manages the fabric data distribution configuration parameters.

Synopsis

```
fddcfg --showall
fddcfg --localaccept policy_list
fddcfg --localreject policy_list
fddcfg --fabwideset policy_list
```

Description Use this command to manage the fabric data distribution configuration parameters. These parameters control the fabric-wide consistency policy.

Switches can be locally configured to allow or reject a security policy. Supported policies include the following:

| | |
|----------|--------------------------------------|
| SCC | Switch Connection Control policy |
| DCC | Device Connection Control policy |
| SCC | Switch Connection Control policy |
| PWD | Password policy |
| FCS | Fabric Configuration Server policy |
| PWD | Fabric Element Authentication policy |
| IPFILTER | IP Filter policy |

Automatic distribution of a fabric-wide consistency policy is limited to SCC, DCC, and FCS policies. Use the **--fabwideset** parameter to enforce these policies fabric-wide in tolerant or strict mode. In strict mode, fabric-wide enforcement of FCS consistency policy is possible in mixed fabrics. However, switches that do not support the policies ignore them.

Notes When a policy is set to strict mode, the manual distribution of the policy is not allowed.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--showall Displays the accept/reject configuration of all policy sets and the fabric-wide consistency policy on the switch.

--localaccept *policy_list*

Configures the switch to accept distributions of the specified policies. The policies in *policy_list* must be separated by semicolons and enclosed in quotation marks; for example, "SCC;DCC;FCS".

--localreject *policy_list*

Configures the switch to reject distributions of the specified policies in *policy_list*. However, a database cannot be rejected if it is specified in the fabric-wide consistency policy. The policies in *policy_list* must be separated by semicolons and enclosed in quotation marks; for example, "SCC;DCC".

--fabwideset *policy_list*

Sets the fabric-wide consistency policy. A database that is set to reject distributions cannot be specified in the fabric-wide consistency policy. To set the fabric-wide consistency policy as strict, use the strictness indicator "S". To set the fabric-wide consistency policy as tolerant, omit the "S". A valid policy set should be of the form "SCC:S;DCC;FCS". To set the fabric-wide policy to NULL (default) or no fabric-wide consistency, use the policy Set "". Supported policies are Switch Connection Control (SCC), Device Connection Control (DCC), and Fabric Configuration Server (FCS). All members specified in a given policy set are automatically distributed to all participating switches in the fabric that support the policy. Refer to the DESCRIPTION section for specific exceptions. In the presence of a fabric-wide FCS consistency policy, this command can only be run from the primary FCS switch.

Examples

To display the fabric-wide consistency policy and the accept/reject configuration for all databases:

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-----
      SCC - accept
      DCC - accept
      PWD - accept
      FCS - accept
      AUTH - accept
      IPFILTER - accept
Fabric Wide Consistency Policy:- "SCC:S;DCC;FCS"
```

To configure the switch to accept distribution of the SCC policy set and PWD database:

```
switch:admin> fddcfg --localaccept "SCC;PWD"
Local Switch Configured to accept policies.
```

To configure this switch to reject distribution of SCC and DCC policy sets:

```
switch:admin> fddcfg --localreject "SCC;DCC"
Local Switch Configured to reject policies.
```

To set the fabric-wide consistency policy to "strict" for SCC and "tolerant" for DCC and FCS:

```
switch:admin> fddcfg --fabwideset "SCC:S;DCC;FCS"
```

SEE ALSO [distribute](#)

fdmiCacheShow

Displays abbreviated remote FDMI device information, according to remote domain ID.

Synopsis `fdmicacheshow`

Description Use this command to display FDMI cache information for remote domains only.

The state of each remote domain, identified by its domain ID, is shown to be unknown, known, unsupported, or error.

The revision of the switch also displays, followed by the world wide name of the switch.

For HBAs, only the HBA identifiers and registered port lists are displayed. No detailed HBA attributes are displayed. For registered ports, only port identifier and corresponding HBA are shown; no detailed port attributes are displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the FDMI cache:

```
switch:admin> fdmicacheshow
Switch entry for domain 3
  state:   known
  version: v310
  wwn:     10:00:00:60:69:90:03:c7

HBAs:
  10:00:00:00:c9:25:9b:96

Ports: 1
  10:00:00:00:c9:25:9b:96

Total count of devices on the switch is 1
```

SEE ALSO [fdmiShow](#)

fdmiShow

Displays detailed FDMI information.

Synopsis `fdmishow [-hexoutput]`

`fdmishow --help`

Description Use this command to display Fabric-Device Management Interface (FDMI) information for all Host Bus Adapters (HBAs) and ports.

Detailed FDMI information is displayed for local HBAs and ports.

Only abbreviated FDMI information is shown for HBA and ports on remote switches.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-hexoutput

Displays raw hex data for all non-ASCII and non-WWN FDMI port attributes. This operand is optional.

--help

Displays the command usage.

Examples To display FDMI information on a local switch:

```
switch:admin> fdmishow
Local HBA database contains:
 10:00:8c:7c:ff:01:eb:00
Ports: 1
 10:00:8c:7c:ff:01:eb:00
  Port attributes:
    FC4 Types: FCP
    Supported Speed: 2 4 8 16 Gb/s
    Port Speed: 16 Gb/s
    Frame Size: 2112 bytes
    Device Name: bfa
    Host Name: X3650050014
    Node Name: 20:00:8c:7c:ff:01:eb:00
    Port Name: 10:00:8c:7c:ff:01:eb:00
    Port Type: N_PORT (0x1)
    Port Symb Name: port2
    Class of Service: 3
    Fabric Name: 10:00:00:05:1e:e5:e8:00
    FC4 Active Type: FCP
    Port State: 0x5
    Discovered Ports: 0x2
    Port Identifier: 0x030200
HBA attributes:
  Node Name: 20:00:8c:7c:ff:01:eb:00
  Manufacturer: Brocade
  Serial Number: BUK0406G041
```

fdmiShow

```

Model: Brocade-1860-2p
Model Description: Brocade-1860-2p
Hardware Version: Rev-A
Driver Version: 3.2.0.0705
Option ROM Version: 3.2.0.0_alpha_bld02_20120831_0705
Firmware Version: 3.2.0.0_alpha_bld02_20120831_0705
OS Name and Version: Windows Server 2008 R2 Standard | N/A
Max CT Payload Length: 0x00000840
Symbolic Name: Brocade-1860-2p | 3.2.0.0705 | X3650050014 |
Number of Ports: 2
Fabric Name: 10:00:00:05:1e:e5:e8:00
Bios Version: 3.2.0.0_alpha_bld02_20120831_0705
Bios State: TRUE
Vendor Identifier: BROCADE
Vendor Info: 0x31000000

```

```

Local Port database contains:
10:00:8c:7c:ff:01:eb:00

```

```

Remote HBA database contains:
10:00:00:05:1e:ea:05:fa
Ports: 1
10:00:00:05:1e:ea:05:fa

```

```

Remote Port database contains:
10:00:00:05:1e:ea:05:fa

```

To display FDMI information in raw hex format (for non-ASCII and non-WWN FDMI port attributes) on a local switch:

```

switch:admin> fdmishow -hexoutput
Local HBA database contains:
10:00:8c:7c:ff:01:eb:00
Ports: 1
10:00:8c:7c:ff:01:eb:00
Port attributes:
FC4 Types: 0x000001000000000000000000000000000000000000000000000000000000000000000000000000000000
Supported Speed: 0x0000003a
Port Speed: 0x00000020
Frame Size: 0x00000840
Device Name: bfa
Host Name: X3650050014
Node Name: 20:00:8c:7c:ff:01:eb:00
Port Name: 10:00:8c:7c:ff:01:eb:00
Port Type: 0x1
Port Symb Name: port2
Class of Service: 0x08000000
Fabric Name: 10:00:00:05:1e:e5:e8:00
FC4 Active Type:
0x000001000000000000000000000000000000000000000000000000000000000000000000000000000000
Port State: 0x00000005
Discovered Ports: 0x00000002
Port Identifier: 0x00030200
HBA attributes:
Node Name: 20:00:8c:7c:ff:01:eb:00
Manufacturer: Brocade
Serial Number: BUK0406G041
Model: Brocade-1860-2p
Model Description: Brocade-1860-2p

```

```
Hardware Version: Rev-A
Driver Version: 3.2.0.0705
Option ROM Version: 3.2.0.0_alpha_bld02_20120831_0705
Firmware Version: 3.2.0.0_alpha_bld02_20120831_0705
OS Name and Version: Windows Server 2008 R2 Standard | N/A
Max CT Payload Length: 0x00000840
Symbolic Name: Brocade-1860-2p | 3.2.0.0705 | X3650050014 |
Number of Ports: 2
Fabric Name: 10:00:00:05:1e:e5:e8:00
Bios Version: 3.2.0.0_alpha_bld02_20120831_0705
Bios State: TRUE
Vendor Identifier: BROCADE
Vendor Info: 0x31000000
```

```
Local Port database contains:
  10:00:8c:7c:ff:01:eb:00
```

```
Remote HBA database contains:
  10:00:00:05:1e:ea:05:fa
Ports: 1
  10:00:00:05:1e:ea:05:fa
```

```
Remote Port database contains:
  10:00:00:05:1e:ea:05:fa
```

SEE ALSO [fdmiCacheShow](#)

femDump

Collects limited debug information to a remote directory through FTP, SFTP, or SCP. This command is a minimal version of **supportsave** command.

Synopsis `femdump [-u user_name -p password -h host_ip -d remote_dir -l protocol]`

Description Use this command to collect limited debug information to a remote directory through FTP or SCP. This command is a minimal version of a **supportsave** command.

The files generated by this command are compressed before being sent off the switch. The core files and panic dumps remain on the switch after the command is run. The FFDC data are removed after the command has finished.

If there are blade processor (BP) blades installed on the switch, a support file (a tar.gz file) is generated from each slot.

This command accepts IPv4 and IPv6 addresses. If the configured IP address is in IPv6 format, the RAS auto-file transfer and event notification to syslog will not work in the case where the Fabric OS version is downgraded. You must reconfigure auto-file transfer and syslog with IPv4 IP addresses.

In a Virtual Fabric environment, the command saves all chassis-based information and iterates through the defined switch-based information for all logical switches. Chassis permissions are required to execute this command.

Note that quotes should be used around path entries to ensure proper handling of special shell characters.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

If you use anonymous FTP to run **femDump** on a chassis with multiple AP blades, configure the FTP Windows server to allow unlimited anonymous users.

Operands This command has the following operands:

-u *user_name*

Specifies the user name for the FTP, SCP, or SFTP server. This operand is optional; if omitted, anonymous FTP is used.

-p *password*

Specifies the password for the FTP, SCP, or SFTP server. This operand is optional with FTP; if omitted, anonymous FTP is used.

-h *host_ip*

Specifies the IPv4 or IPv6 address for the remote server.

-d *remote_dir*

Specifies the remote directory to which the file is to be transferred.

-l *protocol*

Specifies the transfer protocol. Valid values are File Transfer Protocol (FTP), Secure Copy (SCP), or Secure File Transfer Protocol (SFTP).

If you plan to use SCP to transfer files, it is important to test the command prior to its use with various SCP-mode services. Because the **femDump** command makes several access requests to copy files, it is important that the SCP-mode service be configured so that passwords are not required for each attempted transfer. Failure to configure the service correctly may result in significant delays in obtaining transferred output from the **femDump** command.

When using SCP, **femDump** may create a directory if it does not already exist and the parent directory has the appropriate permissions. Use of FTP requires the directory to exist on the remote server.

Examples To collect limited debug information to a remote directory through SCP:

```
switch:admin> femdump -u admin -h 10.70.4.104 -d /temp/support -l scp
Password:
Saving support information for switch:SW_178, module:F_SSHOW_ASICDB...
Saving support information for switch:SW_178, module:F_SSHOW_FABRIC...
Saving support information for switch:SW_178, module:F_SSHOW_SYS...
Saving support information for switch:SW_178, module:F_SSHOW_PORT...
Saving support information for switch:SW_178, module:F_SSHOW_SERVICE...
Saving support information for switch:SW_178, module:F_SSHOW_ISWITCH...
Saving support information for switch:SW_178, module:F_C3REGDUMP...
Saving support information for switch:SW_178, module:F_C2REGDUMP...
Saving support information for switch:SW_178, module:F_C1REGDUMP...
Saving support information for switch:SW_178, module:F_PBREGDUMP...
Saving support information for switch:SW_178, module:F_BLSREGDUMP...

SupportSave completed.
```

SEE ALSO [supportSave](#), [supportShow](#), [supportFtp](#)

ficonCfg

Configures the specified FICON database.

Synopsis `ficoncfg --set database port`

`ficoncfg --reset database`

`ficoncfg --help`

Description Use this command to configure a FICON database on a specified port. Refer to **ficonShow** for a description of the database content.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

port

Specifies the port to be configured. You can specify the port by its port index number or by a hexadecimal number that contains the Area field (middle byte) of the three-byte Fibre Channel Port Address. Use **switchShow** for a listing of valid port index numbers.

--set

Sets the configuration entry.

--reset

Resets the configuration entry to its default value.

database

Specifies the name of the FICON database. The only database currently supported is the following:

LIRR

Devices registered to receive link incident reports.

Examples To set the LIRR database on a port using a decimal index number:

```
switch:user> ficoncfg --set LIRR 27
```

To reset the LIRR

```
switch:user> ficoncfg --reset LIRR
```

To set the LIRR database using a hexadecimal index number that contains the Area field (middle byte) of the three-byte Fibre Channel Port Address:

```
switch:user> ficoncfg --set LIRR 0x1b
switch:user> ficonshow LIRR
The Local LIRR database has 0 entries.

Current LIRR device port number: 27 (0x1b)
```

SEE ALSO [ficonHelp](#), [ficonShow](#)

ficonClear

Clears the records from the specified FICON database.

Synopsis `ficonclear database`

Description Use this command to remove records from the local FICON database. The command effect depends on the specified database.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

database

Specifies the name of the FICON database. The databases include the following:

RLIR

Remove all entries from the link incidents database including implicit link incidents (ILIR).

RNID

Remove all the "not current" entries from the device node identification database (the entries are for devices that were previously connected but are no longer online). Note that "current" entries are not removed from the RNID database.

Examples To clear the RLIR database:

```
switch:user> ficonclear RLIR
successfully clear local RLIR Database.
```

To clear the RNID database:

```
switch:user> ficonclear RNID
successfully clear not current
entries from local RNID Database.
```

SEE ALSO [ficonHelp](#), [ficonShow](#)

ficonCupSet

Sets FICON-CUP parameters for a switch.

Synopsis `ficoncupset fmsmode enable | disable | reset`

`ficoncupset modereg bitname 0 | 1`

`ficoncupset MIHPTO seconds`

`ficoncupset CRP PID CHID`

Description Use this command to set FICON-CUP (Control Unit Port) parameters for a switch. All parameters can be set while the switch is online. Changes made by this command take effect immediately. A reboot is not required.

Use **ficonCupShow** to display current settings.

Notes A FICON License is required to enable FMSMODE and to manage ports with FICON CUP. Without a license, FICON CUP traffic will not be allowed.

When FMSMODE is enabled, port names are truncated to 24 characters to be FICON compliant, and a 24-character limit is imposed on all newly created port names.

FICON Management Server (FMS) mode cannot be enabled if port ID (PID) Format 2 is used. Refer to the *Fabric OS Administrator's Guide* for information on PID formats.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

fmsmode

Configures the FICON Management Server (FMS) mode for the switch. Specify one of the following modes:

enable

Enables the FICON Management Server mode. When you enable the switch, you must wait until all ports have come online before enabling fmsmode.

disable

Disables the FICON Management Server mode.

reset

Resets the CUP to a neutral execution state. This command does not modify any other configuration state.

modereg

Set a bit in the FICON-CUP mode register. The following operands are required:

bitname 0 | 1

Specifies a given bit value to be (1) set or not set (0). Valid values for *bitname* are

POSC

Programmed offline state control

| | |
|------|------------------------------|
| ASM | Active=saved mode |
| DCAM | Director clock alert mode |
| ACP | Alternate control prohibited |
| HCP | Host control prohibited |

MIHPTO

Sets the missing interrupt handler primary timeout (MIHPTO) value for the CUP. The following operand is required:

seconds

Specifies the timeout value in seconds. Provide a decimal value in the range between 15 and 600 seconds. The default timeout value is 180 seconds. If a value greater than 63 seconds is specified, the timeout value is rounded down to the closest value divisible by 10. For example, an MIHPTO timeout value of 86 defaults to 80.

CRP PID CHID

Sets the current reporting path (CRP). The reporting path is a CUP mechanism for sending FRU-failure reports to a FICON logical path via FICON protocol. The logical path between the PID and the CHID must exist and be in operational state for this command to succeed. Use `ficonCupShow` with the LP option to display the logical paths on the switch. The following operands are required:

PID

Specifies the Port identifier, which is a three-byte Fibre Channel Port Address.

CHID

Specifies the Channel ID (CHID). The CHID is the Logical Partition (LPAR) identifier supplied as part of the FICON protocol header. The CHID is a 1-byte value in hexadecimal format. The first nibble indicating the Channel Subsystem identifier (a value between 0 and 3) and the second the LPAR within that CSS (a value between 0 and F).

Examples

To enable FMS mode for the switch:

```
switch:admin> ficoncupset fmsmode enable
fmsmode for the switch is now Enabled
```

To set the ASM bit in the mode register for the switch:

```
switch:admin> ficoncupset modereg ASM 1
Active=Saved Mode bit is set to 1
```

To set the MIHPTO value to 60 seconds:

```
switch:admin> ficoncupset MIHPTO 60
MIHPTO has been changed to 60 seconds
```

To set the current reporting path:

```
switch:admin> ficoncupset CRP 082300 1A
```

To display the current reporting path:

```
switch:admin> ficoncupshow LP
FICON CUP Logical Paths for CUP 0x08FE00
```

| PID | LP | Reporting State | Path |
|--------|----|-----------------|------|
| 082300 | 1A | Oper*** | Curr |
| 082300 | 1B | Oper | |
| 082300 | 1C | Oper | |
| 082400 | 1A | Reset | |
| 082400 | 1B | Reset | |
| 612400 | 1E | Reset | Prim |
| 612400 | 1F | Reset | |

To reset the FMS mode for a switch:

```
switch:admin> ficoncupset fmsmode reset
FMS_001(I) - FMSMODE RESET completed - FMSMODE(Disabled)
```

Conditions prior to reset:

```
All CUP states nominal, no conditions reset
```

```
----- END DISPLAY of FMSMODE RESET -----
```

SEE ALSO [ficonCupShow](#)

ficonCupShow

Displays FICON-CUP parameters for a switch.

Synopsis

```
ficoncupshow fmsmode
ficoncupshow modereg [bitname]
ficoncupshow MIHPTO
ficoncupshow DD_LOG
ficoncupshow diag_info
ficoncupshow hlthchk_log
ficoncupshow LP
```

Description Use this command to display FICON-CUP (Control Unit Port) parameters for a switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

fmsmode

Display the FICON Management Server (FMS) mode for the switch.

modereg

Display the FICON-CUP mode register. If no operand is specified, all mode register bit settings are displayed. If a mode register bit name is specified, then only the value of that bit is displayed. A value of 1 indicates that a given mode register bit is set, and 0 indicates that it is not set.

The following operand is optional:

bitname

Displays the specified mode register bit as either set (1) or not set (0). Valid values for *bitname* include the following:

POSC

Programmed offline state control

UAM

User alert mode

ASM

Active=saved mode

DCAM

Director clock alert mode

ACP

Alternate control prohibited

| | |
|-------------|--|
| HCP | Host control prohibited |
| MIHPTO | Displays the FICON-CUP missing interrupt handler primary timeout (MIHPTO) value in seconds. |
| DD_LOG | Displays the latest Director Diagnostics Log. |
| diag_info | Displays diagnostic information for the logical switch such as whether Diagnostic Interval has been set for CUP Diagnostics (if so, then Statistics Sampling is running), along with additional information about Statistics Sampling by the CUP, the detected CUP Diagnostic capabilities and settings for other switches in the fabric.. |
| hlthchk_log | Displays the HealthCheck Logs for the logical switch. It displays Sense Data returned to the FICON host for Asynchronous Error Reporting events generated by the CUP for Congestion, Latency, or Port Fencing events which have been detected by the switch. |
| LP | Displays the logical paths on the switch. For each entry, the command displays the port identifier (PID), the LPAR identifier (CHID), reporting state (operational or reset (=nonoperational)), and reporting path state (current, primary, or alternate). |

Examples To display the FMS mode for the switch:

```
switch:user> ficoncupshow fmsmode
fmsmode for the switch: Enabled
```

To display the mode register for the switch:

```
switch:user> ficoncupshow modereg
  POSC  UAM  ASM  DCAM  ACP  HCP
  -----
    1    0    1    1    1    0
```

To display the ASM bit in the mode register for the switch:

```
switch:user> ficoncupshow modereg ASM
ASM
---
```

To display the MIHPTO value for the CUP:

```
switch:user> ficoncupshow MIHPTO
MIHPTO for the CUP: 60 seconds
```

ficonCupShow

To display the logical paths for the switch:

```
switch:user> ficoncupshow LP
FICON CUP Logical Paths for CUP 0x08FE00
      LP      Operational  Reporting
PID    CHID    State      Path State
-----
082300 1A      Oper***   Current
082300 1B      Oper
082300 1C      Oper
082400 1A      Reset
082400 1B      Reset
612400 1E      Reset     Primary
612400 1F      Reset
```

SEE ALSO [ficonCupSet](#), [ficonHelp](#)

ficonHelp

Displays a list of FICON support commands.

Synopsis `ficonhelp`

Description Use this command to display a list of FICON support commands with descriptions.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display a list of FICON commands:

```
switch:admin> ficonhelp
ficoncfg      Manage FICON configuration
ficonclear    Clears contents of the specified FICON
               management database
ficoncupset   Sets FICON-CUP parameters for a switch
ficoncupshow  Displays FICON-CUP parameters for a switch
ficonhelp     Displays FICON commands
ficonshow     Displays contents of the specified FICON
               management database
```

SEE ALSO None

ficonShow

Displays the contents of the specified FICON database.

Synopsis `ficonshow database [fabric | table | port port_index]`

Description Use this command to display the contents of a FICON database. The **ficonShow** database operand is the name of the database to display. If the fabric operand is absent, the command displays the members of the named database that are local to the switch on which the command was issued. If the fabric operand is present, it must be entered exactly as shown, and this specifies that all members are displayed, both local and remote.

The following information may be displayed, depending on which database you enter and which operands you use with the command:

| | |
|---|--|
| Domain | Displays the domain ID. |
| Fabric WWN | Displays the fabric WWN. |
| Flag | Indicates if the node is valid, not valid, or not current. Flag values are as follows: |
| 0x00 | Indicates the node ID of the storage port is valid. |
| 0x10 | Indicates the node ID of the channel port is valid. |
| 0x20 | Indicates the node ID of the storage port is not current. |
| 0x30 | Indicates the node ID of the channel port is not current. |
| 0x40 | Indicates the node ID of the storage port for the RNID switch for RLIR is not valid. |
| 0x50 | Indicates the node ID of the channel port is not valid. |
| Fmt | Displays the record-registration format. |
| FRU Failure Description | Indicates the FRU failure type as one of the following: |
| WWN card [<i>unit number</i>] | The WWN card |
| Power Supply [<i>unit number</i>] | The Power Supply card |
| Hardware Slot [<i>unit number</i>] | The Hardware Slot |

| | |
|----------------------------------|---|
| Blower [unit number] | The Blower |
| FRU Part Number | Displays the FRU part number. |
| FRU Serial Number | Displays the FRU serial number. |
| Incident Count | Displays the incident count. This number increases by 1 for each incident within the individual switch. |
| Link Incident Description | Same as Link Incident Type. |
| Link Incident Type | Indicates the link incident type as one of the following: <ul style="list-style-type: none"> • Bit-error-rate threshold exceeded • Loss of signal or synchronization • NOS recognized • Primitive sequence timeout • Invalid primitive sequence for port state |
| Listener PID | Same as PID. |
| Listener Port Type | Same as Port Type. |
| Listener Port WWN | Displays the channel HBA port world wide name. |
| Listener Type | Indicates the listener type as follows: <ul style="list-style-type: none"> Conditional This port receives a link incident record if no other recipients from the established registration list have been chosen. Unconditional This port is always chosen as a recipient of a link incident record. |
| Manufacturer | Displays the manufacturer name or code. |
| Model Number | Displays the model number. |
| Node Parameters | Same as Parameters. |
| Parameters | Displays the node type for the switch in three bytes, 0xAABBCC: <ul style="list-style-type: none"> Byte AA 0x20 FC-SB-2 and updates |

| | |
|---------------------|---|
| Byte BB 0x0a | Switch |
| Byte CC 0x00 | Port number. It is dynamically assigned whenever a link incident occurs. |
| Parm | Displays the incident node parameters type in three bytes, 0xAABBCC: |
| Byte AA | Possible values include the following: |
| 0x00 | Reserved. |
| 0x20 | FC-SB-2 and updates. |
| 0x40 | Other FC-4s including FCP and updates. |
| 0x60 | FC-SB-2 and updates and other FC-4s including FCP and updates. |
| 0x80 | FC-4 support not specified. |
| 0xa0 | Reserved. |
| 0xc0 | Reserved. |
| 0xe0 | Vendor-specific. |
| Byte BB | Possible values include the following: |
| 0x00 | Unspecified class. |
| 0x01 | Direct access storage device, if it is a storage port; otherwise, not channel-to-channel capable. |
| 0x02 | Magnetic tape, if it is a storage port; otherwise, a reserved field for a channel port. |
| 0x03 | Input unit record, if it is a storage port; otherwise, a reserved field for a channel port. |
| 0x04 | Output unit, if it is a storage port; otherwise, a reserved field for a channel port. |
| 0x05 | Reserved field for a channel port. |

| | |
|-----------------------------|--|
| 0x06 | Controller, if it is a storage port; otherwise, a reserved field for a channel port. |
| 0x07 | Terminal - Full screen if it is a storage port; otherwise, a reserved field for a channel port. |
| 0x08 | Terminal - Line mode if it is a storage port; otherwise, an emulated control unit support only. |
| 0x09 | Reserved. |
| 0x10 | Switch, if it is a switch device; otherwise, reserved. |
| 0x0b-0xff | Reserved. |
| Byte CC | Possible values include the following: |
| 0x00 | If storage CU port has registered with the switch. |
| 0xID | CHIPID if channel port has registered with the switch. |
| 0xPN | If switch has registered with the channel, PN represents the FL port number. |
| Part Number | Displays the switch chassis part number. |
| PID | Displays the 24-bit Fibre Channel port address in OxDDAAPP format. DD is Domain ID. AA is Area ID. PP is AL_PA ID. |
| Plant of Manufacture | Displays the manufacturer plant name or code. |
| Port | Physical port number. |
| Port Status | Displays the status of the port as one of the following: <ul style="list-style-type: none"> • Link degraded but operational • Link not operational |
| Port Type | Displays the port type as one of the following: |
| U | Unknown |

| | |
|----------------------------|--|
| N | N_Port |
| NL | NL_Port |
| Protocol | Displays whether the traffic is using FICON or FCP. |
| Registered Node WWN | Displays the device's node world wide name associated with the device HBA. |
| Registered Port WWN | Displays the device's channel or storage CU port world wide name associated with the device HBA. |
| Sequence Number | Displays the sequence number of the self-describing node. |
| Serial Number | Displays the switch serial number. |
| Switch node WWN | Displays the switch node world wide name. |
| Switch Port WWN | Displays the switch port world wide name. |
| Switch WWN | Displays the switch WWN. |
| Tag | Displays the physical identifier for the self-describing node interface. |
| TS Format | Displays the Time Server format. |
| Time Stamp | Displays the timestamp, expressed in date format. |
| Type | Same as Port Type. |
| Type Number | Displays the type number of the self-describing node. It also describes the machine type. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------------|---|
| <i>database</i> | Specify the database to display. This operand is required. Valid values include the following |
| RNID | Device node identification. |

| | |
|-------------------|--|
| LIRR | Devices registered to receive link incident reports. Entries can have the following flags: |
| -C* | A user-configured LIRR entry defined as current (set with the ficonCfg command). |
| -C | A LIRR entry defined by the management server demon as current based on order. |
| -S | A LIRR entry defined by the management server demon as secondary. |
| SwitchRNID | Switch node identification. |
| RLIR | Link incident reports. |
| ILIR | Implicit link incident reports. |
| fabric | Displays FICON database information for the entire fabric. This operand is optional; if omitted, only local members of the named database are displayed. |
| table | Displays RNID data in table format. |
| port | Displays RNID data for the specified port. |
| port_index | Specifies the port index number in decimal or the area field (middle byte) of the Fibre Channel Address (PID) in hexadecimal format. |

Examples To display the local RNID database:

```
switch:admin> ficonshow RNID
{
  {Fmt Type PID Registered Port WWN      Registered Node WWN \
0x18 N 502b00 50:05:07:64:01:00:15:8d 50:05:07:64:00:c1:69:ca \
  flag Parm
  0x10 0x200110
  Type number:      002064
  Model number:     101
  Manufacturer:     IBM
  Plant of Manufacture: 02
  Sequence Number:  0000000169CA
  tag:              102b
}
{Fmt Type PID Registered Port WWN      Registered Node WWN \
0x18 N 502e00 50:05:07:64:01:40:0f:ca 50:05:07:64:00:c1:69:ca \
  flag Parm
  0x10 0x200105
  Type number:      002064
  Model number:     101
  Manufacturer:     IBM
```

```

Plant of Manufacture: 02
Sequence Number:      0000000169CA
tag:                  052e
}
}

```

To display the local RNID database in tabular format:

```

switch:admin> ficonshow RNID table
{
Fmt  Type PID      Registered Port WWN      Registered Node WWN
0x18 N   252500 50:05:07:60:28:bf:42:cf 50:05:07:64:00:cd:01:b6
0x18 N   255800 50:05:07:60:28:bf:3e:98 50:05:07:64:00:cd:01:b6
0x18 N   255A00 50:05:07:60:28:bf:2c:9e 50:05:07:64:00:cd:01:b6
0x18 N   255B00 50:05:07:60:28:bf:3f:0b 50:05:07:64:00:cd:01:b6
0x18 N   255C00 50:05:07:60:28:bf:3f:0a 50:05:07:64:00:cd:01:b6
0x18 N   255D00 50:05:07:60:28:bf:2c:9f 50:05:07:64:00:cd:01:b6
0x18 N   255E00 50:05:07:60:28:bf:18:63 50:05:07:64:00:cd:01:b6
0x18 N   255F00 50:05:07:60:28:bf:18:64 50:05:07:64:00:cd:01:b6
0x18 E   258600 20:86:00:05:33:0d:b7:05 10:00:00:05:33:0d:b7:05
0x18 E   258700 20:87:00:05:33:0d:b7:05 10:00:00:05:33:0d:b7:05
0x18 E   25F600 20:f6:00:05:33:0d:b7:05 10:00:00:05:33:0d:b7:05
      flag Parm      Type  Mod  Manf  Plant  Sequence#  Tag
      0x10 0x3101DB 002817 M15  IBM   02      0000000D01B6 40DB
      0x10 0x310138 002817 M15  IBM   02      0000000D01B6 4038
      0x10 0x3101D4 002817 M15  IBM   02      0000000D01B6 40D4
      0x10 0x3101D7 002817 M15  IBM   02      0000000D01B6 40D7
      0x10 0x3101D6 002817 M15  IBM   02      0000000D01B6 40D6
      0x10 0x3101D5 002817 M15  IBM   02      0000000D01B6 40D5
      0x10 0x3101D8 002817 M15  IBM   02      0000000D01B6 40D8
      0x10 0x3101D9 002817 M15  IBM   02      0000000D01B6 40D9
      0x00 0x200A00 SLKWRM DCX  BRD   CA      5FT00X39C19A 22FF
      0x00 0x200A00 SLKWRM DCX  BRD   CA      5FT00X39C19A 22FF
      0x00 0x200A00 SLKWRM DCX  BRD   CA      5FT00X39C19A 22FF
}
11 valid entries, 0 not current entries
The Local RNID database has 11 entries.

```

To display RNID data for the specified port:

```

switch:admin> ficonshow RNID port Ox08
{Fmt  Type PID      Registered Port WWN      Registered Node WWN  \
0x18 E   010800 20:08:00:05:1e:57:b1:86 10:00:00:05:1e:57:b1:86\
      flag Parm
      0x00 0x200a00

Type number:      BROCAD
Model number:     510
Manufacturer:     BRD
Plant of Manufacture: CA
Sequence Number:  0ALM0632D038
tag:              03ff
}

```

To display the local LIRR database:

```
switch:admin> ficonshow LIRR
{Fmt  Type PID      Listener Port WWN
 0x18 N    255800 50:05:07:60:28:bf:3e:98 \
 0x18 N    255a00 50:05:07:60:28:bf:2c:9e \
 0x18 N    255b00 50:05:07:60:28:bf:3f:0b \
 0x18 N    255c00 50:05:07:60:28:bf:3f:0a \
 0x18 N    255d00 50:05:07:60:28:bf:2c:9f \
 0x18 N    255e00 50:05:07:60:28:bf:18:63 \
 0x18 N    255f00 50:05:07:60:28:bf:18:64 \
```

```
Switch Port WWN      Listener Type
20:58:00:05:33:0d:b7:05 Conditional-S
20:5a:00:05:33:0d:b7:05 Conditional
20:5b:00:05:33:0d:b7:05 Conditional
20:5c:00:05:33:0d:b7:05 Conditional
20:5d:00:05:33:0d:b7:05 Conditional
20:5e:00:05:33:0d:b7:05 Conditional
20:5f:00:05:33:0d:b7:05 Conditional-C*
}
```

The Local LIRR database has 7 entries.

Current LIRR device port number: 95 (0x5f)

To display the local and remote LIRR database:

```
switch:admin> ficonshow LIRR fabric
{Fmt  Type PID      Listener Port WWN      \
 0x18 N    502d00 50:05:07:64:01:40:11:79 \
 0x18 N    510d00 50:05:07:64:01:00:15:8c \
 0x18 N    510f00 50:05:07:64:01:00:14:62 \
```

```
Switch Port WWN      Listener Type
20:2d:00:60:69:80:1e:4e Conditional-C
20:0d:00:60:69:80:1e:4f Conditional-S
 20:0f:00:60:69:80:1e:4f Conditional
}
```

The LIRR database has 3 entries.

Current LIRR device port number: Not configured

To display the local Switch RNID database:

```
switch:admin> ficonshow switchrnid
{
 {Switch WWN          flag Parm
 10:00:00:60:69:80:1e:4e 0x00 0x200a00
 Type number:          SLKWRM
 Model number:         48K
 Manufacturer:         BRD
 Plant of Manufacture: CA
 Sequence Number:      0RB030000082
 tag:                  00ff
 }
}
```

The Local switch RNID database has 1 entries.

ficonShow

To display the local RLIR database:

```
switch:user> ficonshow RLIR

{
  {Fmt Type PID Port Incident Count TS Format Time Stamp
   0x18 N 502e00 46 1 Time server Mon Jan 13 04:29:33 2003
   Port Status: Link not operational
   Link Failure Type: Loss of signal or synchronization

   Registered Port WWN Registered Node WWN Flag Node Parameters
   50:05:07:64:01:40:0f:ca 50:05:07:64:00:c1:69:ca 0x50 0x200105
   Type Number: 002064
   Model Number: 101
   Manufacturer: IBM
   Plant of Manufacture: 02
   Sequence Number: 0000000169CA
   tag: 2e00

   Switch Port WWN Switch Node WWN Flag Node Parameters
   20:2e:00:60:69:80:1e:4e 10:00:00:60:69:80:1e:4e 0x00 0x200a2e
   Switch Part Number: 060-0001501-05
   Switch Serial Number: 0FT02X801E4E
   Domain: 20480
  }
}
The local RLIR database has 1 entry.
```

SEE ALSO [ficonClear](#)

fipsCfg

Configures FIPS (Federal Information Processing Standards) mode.

Synopsis `fipscfg --enable [fips | selftests | bootprom | dh | sha256]
[-nowarn] [-dp] [enable | disable]`

`fipscfg --disable [selftests | bootprom | dh | sha256]
[-nowarn] [-dp] [enable | disable]`

`fipscfg --zeroize [-nowarn] [-dp]`

`fipscfg --show | --showall`

`fipscfg --force fips [-dp]`

`fipscfg --verify fips [-dp]`

Description Use this command to configure FIPS mode on the switch. In this mode, only FIPS-compliant algorithms are allowed. As part of FIPS 140-2 level-2 compliance, passwords, shared secrets and the private keys used in SSL/TLS, system login, etc., need to be zeroized. Power-up self tests are executed when the switch is powered on to check for the consistency of the algorithms implemented on the switch.

This command prompts for confirmation before FIPS configuration changes take effect. Specifying no cancels the operation. The **-nowarn** option overrides the prompting.

Notes Certain services and functions, such as FTP, HTTP, remote procedure calls (RPC), root account, boot prom access, etc., must be blocked before the system can enter FIPS mode.

LDAP should not be configured while FIPS is enabled.

The system must be rebooted for FIPS mode changes to take effect.

Refer to the *Fabric OS Administrator's Guide* for information on configuring your system for FIPS 140-2 level-2 compliance.

FIPS mode cannot be modified through **configDownload**.

FIPS is not supported on all platforms. For FIPS-compliant hardware, refer to the *Fabric OS Administrator's Guide*.

In a Virtual Fabric environment, FIPS is treated as chassis-wide configuration and applies to all logical switches in the chassis. Chassis permissions are required to configure FIPS.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--help
Prints command usage.

--disable selftests [-nowarn]
Disables selftests mode.

- enable [fips | selftests] [-nowarn]**
Enables FIPS or selftests mode. Selftests must be enabled before FIPS mode is enabled.
- zeroize [-nowarn] [-dp]**
Erases all passwords, shared secrets, private keys, etc. in the system. Specify the **-dp** option to zeroize both CP and dual-CP.
- show | --showall**
Displays the current FIPS configuration.
- force fips [-dp]**
This option enables FIPS mode even if prerequisites are not met, except under the following two conditions:
- In a dual-CP system if HA is not in sync between the two CPs.
 - If selftests is in a disabled state.
- Specify the **-dp** option to enable FIPS mode on both CP and dual-CP.
- verify fips [-dp]**
Scans the prerequisites for enabling FIPS and print the failure/success cases. Specify the **-dp** option to verify both CP and dual-CP.
- disable | --enable bootprom [-nowarn]**
Disables or enables the Boot Programmable Read-Only Memory (Boot PROM) on the switch. Boot PROM access is blocked in FIPS mode. Disabling Boot PROM requires root permission. Enabling Boot PROM does not require root permission.
- disable | --enable dh [-nowarn]**
Removes or configures all the Diffie hellman specific SSH configurations on both the SSH server and client configuration, and restarts the SSH server. This command is supported only in FIPS mode
- disable | --enable sha256 [-nowarn]**
Configures the hash as SHA1 (for disable command) or SHA256 for signature generation and verification in both the SSH server and client configuration, and restarts the SSH server. This command is supported only in FIPS mode
- disable | --enable selftests [-dp] [enable | disable]**
Disables or enables the selftests mode on dual-CP or both dual-CP and CP. Specify the **-dp enable** option to enable or disable selftests on both CP and dual-CP. Specify the **-dp disable** option to enable or disable selftests on DP only. Ignore the **-dp** option to enable or disable selftests on the CP only.

Examples

To display the current FIPS configuration:

```
switch:admin> fipscfg --show
FIPS mode is :      Disabled
FIPS Selftests mode/status is :  Disabled/None
```

To enable selftests:

```
switch admin> fipscfg --enable selftests
You are enabling selftests.
Do you want to continue? (yes, y, no, n) [no] : yes
FIPS Selftests mode/status has been set to :      Enabled/None
```

To verify FIPS prerequisites:

```
switch:admin> fipscfg --verify fips
Standby firmware supports FIPS - PASS
SELF tests check has passed - PASS
Root account is enabled - FAIL
Radius check has passed - PASS
Authentication check has passed - PASS
Inflight Encryption check has passed - PASS
IPSec check has passed - PASS
Telnet port number <23> for the policy <default_ipv4> is in permit state.
HTTP port number <80> for the policy <default_ipv4> is in permit state.
RPC port number <897> for the policy <default_ipv4> is in permit state.
IPv4 policies not FIPS compliant - FAIL  NEWLY Added message
Telnet port number <23> for the policy <default_ipv6> is in permit state.
HTTP port number <80> for the policy <default_ipv6> is in permit state.
RPC port number <897> for the policy <default_ipv6> is in permit state.
IPv6 policies not FIPS compliant - FAIL
SNMP is not in read only mode - FAIL
Bootprom access is enabled - FAIL
Firmwaredownload signature verification is disabled - FAIL
Secure config upload/download is disabled - FAIL
SSH DSA Keys check passed - PASS
Inband Management interface is disabled - PASS
Ipseconfig is disabled. - PASS
Signature Algorithm used in LDAP CA certificate is not Sha256
Signature Algorithm used in Switch CA certificate is not Sha256
Signature Algorithm used in HTTPS certificate is not Sha256
Signature Algorithm used in FCAP certificate is not Sha256
Signature Algorithm used in FCAP CA certificate is not Sha256
Certificates are not FIPS compliant - FAIL
Certificates validation has passed - PASS
SSH client's ShaValue is not configured as 2 - FAIL
SSH client's MinPrime is not configured as 2048 - FAIL
SSH client's KexAlgorithms is not configured as diffie-hellman-group-exchange-sha256
- FAIL
SSH client's ServerHostKeySize is not configured as 2048 - FAIL
SSH server's ShaValue is not configured as 2 - FAIL
SSH config is not FIPS compliant - FAIL
```

To enable FIPS after prerequisites have been met:

```
switch:admin> fipscfg --enable fips
You are enabling FIPS.
Do you want to continue? (yes, y, no, n) [no] : yes
FIPS mode has been set to : Enabled
Please reboot the system

switch:admin> fipscfg --show
FIPS mode is : Enabled
```

To attempt enabling FIPS when prerequisites are not met:

```
switch:admin> fipscfg --enable fips
SelfTests mode is not enabled.
Root account is enabled.
Authentication uses MD5 hash algorithm.
Authentication uses DH group 0.
```

```

Telnet port number <23> for the policy <default_ipv4> \
  is in permit state.
HTTP port number <80> for the policy <default_ipv4> \
  is in permit state.
RPC port number <898> for the policy <default_ipv4> \
  is in permit state.
Telnet port number <23> for the policy <default_ipv6>\
  is in permit state.
HTTP port number <80> for the policy <default_ipv6> \
  is in permit state.
RPC port number <898> for the policy <default_ipv6> \
  is in permit state.
SNMP is not in read only mode.
Bootprom access is enabled.

```

FIPS mode cannot be configured at this time

To back out of a zeroizing operation:

```

switch:admin> fipscfg --zeroize
You are Zeroizing FIPS configuration.
Do you want to continue? (yes, y, no, n) [no]: no
Operation cancelled.

```

```

switch:admin> fipscfg --zeroize
You are Zeroizing FIPS configuration.
Do you want to continue? (yes, y, no, n) [no]: yes

```

Executing 'secauthsecret --remove --all':

```

This command deletes database of DH-CHAP secret keys.\
  If a fabric requires authentication, deleting this \
  database may cause switch to segment from the fabric.

```

```

Do want to remove secret key database? \
  (yes, y, no, n): [no]
Operation cancelled...
Executing 'pkiremove':

```

WARNING!!!

Removing pki objects will impair the security functionality of this fibre channel switch. If you want secure mode enabled, you will need to get the switch certificate again.

```

About to remove Pki objects.
ARE YOU SURE (yes, y, no, n): [no]
Operation cancelled.
Executing 'passwddefault':
Password policies are already set to default.
Executing 'seccertutil delkey':

```

Deleting the key pair will automatically do the following:

1. Delete all existing CSRs.
2. Delete all existing certificates.
3. Reset the certificate filename to none.
4. Disable secure protocols.

```

Continue (yes, y, no, n): [no]

```

```
Operation cancelled.  
Zeroizing Radius configuration:  
RADIUS configuration does not exist.  
LDAP configuration does not exist.  
Zeroizing IPsec static SA configuration.  
Zeroizing SSH key.  
Permission denied to perform this operation.  
Failed to zeroize SSH public key.  
switch:admin>
```

To disable bootprom:

```
switch:root> fipscfg --disable bootprom  
You are disabling bootprom.  
Do you want to continue? (yes, y, no, n) [no] : no  
Operation cancelled.
```

SEE ALSO None

firmwareActivate

Activates switch firmware.

Synopsis `firmwareactivate`

`firmwareactivate [-l | -local]`

`firmwareactivate --help`

Description Use this command to activate the firmware that has been downloaded to the secondary partition using the `firmwareDownload -r` command. This command swaps partitions and reboots the system to activate the new image. Execute this command with the `-local` or `-l` option to activate the firmware on a local control processor (CP).

The `firmwareRestore` and `firmwareCommit` commands only take action if the new firmware is activated; otherwise, the process terminates with an error. However, you can execute the `firmwareDownload` command before activating the firmware.

Notes This command is supported on the single CP and dual-CP systems.

You must not make any configuration changes before activating the firmware.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operands are optional:

`-l | -local`
Activates the firmware on the local CP.

`--help`
Displays the command usage.

Examples To activate a new version of the firmware:

```
switch:admin> firmwareactivate
This command will activate the firmware on the secondary partition
but will require that existing telnet, secure telnet or SSH sessions
to be restarted.
```

```
Do you want to continue (Y/N) [Y]:
```

SEE ALSO [firmwareCommit](#), [firmwareCleanInstall](#), [firmwareDownload](#), [firmwareDownloadStatus](#), [firmwareKeyShow](#), [firmwareRestore](#), [firmwareShow](#), [version](#)

firmwareCheck

Performs firmware integrity test.

Synopsis `firmwarecheck`

Description Use this command to check the integrity of software image that is stored on a switch. This command also detects the files whose checksum has changed. The firmware integrity test fails if the files or packages are modified manually.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To perform firmware integrity test:

```
switch:user> firmwarecheck  
Validating integrity of firmware on root filesystem  
Please wait...  
Firmware integrity test passed.
```

SEE ALSO [firmwareCommit](#), [firmwareDownload](#), [firmwareKeyShow](#), [firmwareRestore](#), [firmwareShow](#), [version](#)

firmwareCleanInstall

Recovers the switch firmware.

Synopsis `firmwarecleaninstall`

`firmwarecleaninstall [-p | -protocol protocol] [host,user,path,password]`

`firmwarecleaninstall --help`

Description Use this command to initiate a clean reinstall of the firmware in cases where the loaded firmware does not function correctly, the normal firmware download fails, or to recover from a rolling reboot situation.

Notes The `firmwarecleaninstall` operation should be performed similarly on both standalone platforms and on directors. On a director, it should be invoked on each CP, and operates only on that CP. `firmwarecleaninstall` can be run on both active and standby CPs, however it should not be run simultaneously on both CPs.

If for any reason, boot flash is corrupted and boot ROM cannot be loaded, you must let it go through RMA. `firmwarecleaninstall` mechanism cannot be used to recover from such damage.

If local CF is corrupted in a way that the Linux image cannot be loaded, `firmwarecleaninstall` cannot be used to recover the switch.

Since pre-existing configurations are totally wiped out by `firmwarecleaninstall`, to restore the previous configuration or the OEM configuration you must do a specific `configdownload`.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operands are optional. When invoked without operands, the command goes into interactive mode.

`-p | -protocol protocol`

Specifies the file transfer protocol. Valid values are `ftp`, `scp`, and `sftp`.

`host`

Specifies a valid FTP or SSH server name or IP address. The firmware is downloaded from the specified host.

`user`

Specifies a user name for FTP or SSH server access.

`path`

Specifies a path for the firmware files.

`passwd`

Specifies a password.

`--help`

Displays the command usage.

Examples To perform a clean firmware install interactively:

```
switch:admin> firmwarecleaninstall
Server Name or IP Address: 10.10.2.10
User Name: admin
File Name: /admin/dist/FOS7.4.0/
Network Protocol(1-FTP, 2-SCP, 3-SFTP) [1]: 1
Password:
```

This command will erase all the data on the Compact Flash before installing a new firmware. There are also disruptive reboots during the process. Please use "configupload -all" if all configurations need to be saved before continuing with firmwarecleaninstall.

```
Do you want to continue (Y/N) [Y]:
```

To perform a clean firmware install non-interactively:

```
switch:admin> firmwarecleaninstall 10.10.2.10 \
admin,/admin/dist/FOS7.4.0/
```

SEE ALSO [firmwareActivate](#), [firmwareCommit](#), [firmwareDownload](#), [firmwareDownloadStatus](#), [firmwareKeyShow](#), [firmwareRestore](#), [firmwareShow](#), [version](#)

firmwareCommit

Commits switch firmware.

Synopsis `firmwarecommit`

Description Use this command to commit a firmware download to a CP. This command copies an updated firmware image to the secondary partition and commits both partitions of the CP to an updated version of the firmware. This must be done after each firmware download and after the switch has been rebooted and a sanity check is performed to make sure the new image is fine.

For switches that have nonvolatile memory set into two equal partitions, the primary partition is the where the system boots from; the secondary partition is where a copy of the firmware is stored, in case the primary partition is damaged.

To maintain the integrity of the firmware image in the nonvolatile memory, the **firmwareDownload** command updates the secondary partition only. When **firmwareDownload** completes successfully and the CP is rebooted, the system switches the primary partition (with the old firmware) to the secondary, and the secondary partition (with the new firmware) to the primary.

The default behavior of the **firmwareDownload** command is to automatically run the **firmwareCommit** command after the reboot. If you decide to disable the autocommit option when running **firmwareDownload**, you must execute one of the following two commands after the CP is rebooted:

- **firmwareCommit** copies the primary partition (with new firmware) to the secondary and commits the new firmware to both partitions of the CP.
- **firmwareRestore** copies the secondary partition (with the old firmware) to the primary and backs out of the new firmware download. The **firmwareRestore** command can be run only if autocommit was disabled during the firmware download. Autocommit can be disabled only when you run **firmwareDownload** in single mode.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To commit a new version of the firmware:

```
switch:admin> firmwarecommit
Validating primary partition...
Doing firmwarecommit now.
Please wait ...
Replicating kernel image
.....
FirmwareCommit completes successfully.
```

SEE ALSO [firmwareDownload](#), [firmwareRestore](#)

firmwareDownload

Downloads firmware from a remote host, a local directory, or a USB device.

Synopsis To invoke the command in interactive mode:

```
firmwaredownload
```

To download FOS firmware over a network:

```
firmwaredownload [-s [-b | -n]] [-p ftp | scp | sftp]
                  [-A cra] [-c] [-r] [-lr] [host, user, path, password]
```

To download FOS firmware from a USB device:

```
firmwaredownload [-s [-b | -n]] [-U]
                  [-c] [-r] [-lr] path
```

Description

Use this command to download switch firmware from an FTP or SSH server or local NFS directory to nonvolatile storage. Switch firmware can also be downloaded from an external USB device on platforms that support USB.

The new firmware is downloaded as a bundle of packages. Package names are defined in a *.plist file along with other firmware information (time stamp, platform code, version, etc.). These packages are made available periodically to add features or to remedy defects. Contact customer support to obtain information about available firmware versions.

On enterprise-class platforms, this command, by default, downloads the firmware image to both control processors (CPs) in rollover mode to prevent disruption to application services. This operation depends on High Availability (HA) support. If HA is not available, use the **-s** option to upgrade the CPs one at a time.

All systems supported by this firmware have two partitions of nonvolatile storage (primary and secondary) to store two firmware images. This command always downloads the new image to the secondary partition and then swaps partitions so the secondary partition becomes the primary.

By default, **firmwareDownload** reboots the system and activates the new image. Finally, the command performs a **firmwareCommit** automatically to copy the new image to the other partition. In systems with blade processors (BPs), after the new CP firmware is downloaded to the system and activated, the BP firmware is downloaded to the BP processors if there is a mismatch between the BP and CP firmware.

By default, **firmwareDownload** performs a full install, autoreboot, and autocommit. These modes are selectable only in single CP (**-s**) mode, in which case autoreboot is OFF by default.

For each standalone switch in your fabric, complete all firmware download changes before issuing the **firmwareDownload** command on the next switch to ensure a nondisruptive download.

If **firmwareDownload** is interrupted due to an unexpected reboot as a result of a software error or power failure, the command automatically recovers the corrupted secondary partition. Wait for the recovery to complete before starting another firmware download.

Notes Firmware download and subsequent POST failure may occur on some platforms, if the firmware commit operation coincides with the execution of POST. The recommended work around is to disable POST (**diagDisablePost**) before you initiate a firmware download and re-enable POST (**diagEnablePost**) after the firmware commit operation has completed.

Firmware download procedures may vary depending on which Fabric OS version you are migrating from. See the *Fabric OS Administrator's Guide* for restrictions on changing Fabric OS versions.

On certain occasions, you may see messages in the console output of **firmwareDownload**. These are internal messages generated by the Linux utilities. You can safely ignore these messages. For relevant Brocade-generated firmware download messages, refer to the *Fabric OS Message Reference* (SULB module).

To correlate Brocade blade names with blade IDs, use the **slotShow** command.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operands are optional. When invoked without operands, the command goes into interactive mode.

-U

Downloads the firmware from an attached USB device. This option is valid only on platforms that support a USB port. Refer to your specific Hardware Reference Guide for details. The USB device must be enabled prior to firmware download with the **usbStorage** command. Firmware must be stored under the default firmware directory in the USB file system. On a dual-CP chassis, the USB device must be attached to the active CP. When downloading firmware from a USB device, the **-p** option is ignored.

-s

Enables single-CP mode. This mode supports selectively enabling or disabling a full install, autoreboot, and autocommit on bladed and nonbladed systems. On enterprise-class platforms, this mode supports upgrading a single CP. When downloading the main Fabric OS firmware, this option disables autoreboot, unless overridden by the **-b** option.

-b

Enables autoreboot mode. When single CP mode is enabled and this operand is not specified, **reboot** must be run manually to activate the downloaded image. If autoreboot mode is enabled, the switch reboots automatically after the firmware has been downloaded.

-n

Disables autocommit mode. When autocommit mode is disabled, the **firmwareCommit** command must be executed manually to propagate the downloaded image to both partitions of the storage device.

host

Specify a valid FTP or SSH server name or IP address. IPV4 and IPV6 addresses are supported. The firmware is downloaded from the specified host. If a host is not specified, the firmware is considered accessible on a local directory. To mention an FTP server by name, a DNS server must first be set up with the **dnsConfig** command. If DNS is enabled and a server name is specified, **firmwareDownload** automatically determines whether IPV4 or IPV6 should be used.

user

Specify a user name for FTP or SSH server access. This operand can be omitted, if the firmware is accessible on a local directory, a USB device, or by anonymous FTP server access. A user name other than "anonymous" is required for SSH server access.

path

Specify a fully qualified path for the firmware. Absolute path names may be specified using forward slashes (/).

password

Specify a password. This operand can be omitted, if the firmware is accessible through a local directory or an attached USB device, or if no password is required by the FTP server. This operand is required when accessing an SSH server. In Fabric OS v7.0.0 or later, you can omit the password if the switch is configured as follows:

- The switch must be configured with public key authentication. Refer to the **sshUtil** help page for more information.
- You select a secure protocol (SCP or SFTP).
- The private key is installed on the switch, and the public key is exported to the remote host. Refer to the *Fabric OS Administrator's Guide* for configuration procedures.

-p scp | ftp | sftp

Specify the file transfer protocol. Valid values are **ftp** (file transfer protocol), **sftp** (secure file transfer protocol), and **scp** (secure copy protocol). Values are not case-sensitive. If **-p** is not specified, **firmwareCommit** determines the protocol automatically by checking the `config.security` parameter.

-A cra

Specify the method for protocol. Valid option is **cra**. Challenge Response Authentication (CRA) is supported only with the SCP protocol.

-c

Disables version compatibility checking. By default, **firmwareDownload** checks if the firmware being downloaded is compatible with other running firmware images in the system. If the firmware version is not compatible, **firmwareDownload** fails. If this option is specified, version compatibility checking is disabled.

-r

Downloads the firmware to the secondary partition only.

-lr

Downloads the firmware locally to the secondary partition on the CP.

Diagnostics The command checks the network connection and other system parameters before initiating **firmwareDownload**. It may fail if at least one of the following conditions is encountered:

- The host is not reachable from the switch.
- The user does not have permission on the host.
- The password is not specified correctly.
- Indicated firmware does not exist on the host, or is not in the right format, or is corrupted.
- The FTP or SSH service is not running on the host.
- The platform is not supported by the firmware indicated.
- The USB device may not be plugged in correctly. On standalone switches, the device must be plugged into the switch USB port. On enterprise-class platforms, the USB device must be plugged into the Active CP
- The USB device is not enabled. Use the **usbStorage** command on the switch to enable the USB device. On enterprise-class platforms, the command must be run on the Active CP to enable the USB device.

For additional system messages, refer to the *Fabric OS Message Reference*.

Examples To download the firmware to both CPs on a dual-CP chassis with an attached USB device (You would execute the same command on a single-CP switch with USB support. Output may vary depending on platform.):

```
switch:admin> firmwaredownload -U v7.0.0

Checking system settings for firmwaredownload...
Protocol selected: USB
Trying address-->AF_INET IP: 127.1.1.8, flags : 2
System settings check passed.

Checking version compatibility...
Version compatibility check passed.

This command will upgrade the firmware on both
CP blades. If you want to upgrade firmware on a single
CP only, please use -s option.

You may run firmwaredownloadstatus to get the status
of this command.

This command will cause a warm/non-disruptive boot on
the active CP, but will require that existing telnet,
secure telnet or SSH sessions be restarted.
```

To download firmware interactively:

```
switch:admin> firmwaredownload
Server Name or IP Address: 192.168.32.10
User Name: admin
File Name: ~admin/dist/FOS7.0.0/
Network Protocol(1-auto-select, 2-FTP, 3-SCP, 4-SFTP) [1]:
Password:*****
Do you want to continue with CRA (Y/N) [N]:
```

To download firmware without a password using public key authentication:

```
switch:admin> firmwaredownload
Server Name or IP Address: 10.38.17.17
User Name: admin
File Name: /users/home25/admin/dist
Network Protocol(1-auto-select, 2-FTP, 3-SCP, 4-SFTP) [1]: 3
Verifying if the public key authentication is available.
Please wait ...Success.

Do Auto-Commit after Reboot [Y]: y
Reboot system after download [N]: y
Server IP: 10.38.17.17, Protocol IPv4
Checking system settings for firmwaredownload...
System settings check passed.
(Output truncated)
```

SEE ALSO [firmwareCommit](#), [firmwareDownloadStatus](#), [firmwareKeyShow](#), [firmwareRestore](#), [firmwareShow](#), [reboot](#), [slotShow](#), [version](#)

firmwareDownloadStatus

Displays the status of a firmware download.

Synopsis `firmwaredownloadstatus`

Description Use this command to display an event log that records the progress and status of events during a firmware download. The event log is created by the firmware download process and is kept until you issue another **firmwareDownload** command. A timestamp is associated with each event. When downloading Fabric OS firmware, the event logs in the two control processors (CPs) are synchronized and you can run this command from either CP.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the status of a firmware download on a switch:

```
switch:admin> firmwaredownloadstatus
[1]: Tue Mar 29 21:43:14 2011
Firmware is being downloaded to the switch. This step may take up to 30 minutes.

[2]: Tue Mar 29 21:48:26 2011
Firmware has been downloaded to the secondary partition of the switch.

[3]: Tue Mar 29 22:06:57 2011
The firmware commit operation has started. This may take up to 10 minutes.

[4]: Tue Mar 29 22:10:19 2011
The commit operation has completed successfully.

[5]: Tue Mar 29 22:10:19 2011
Firmwaredownload command has completed successfully. Use firmwaredownloadstatus to verify the
firmware versions.
```

To display the status of a firmware download on a chassis:

```
switch:admin> firmwaredownloadstatus
[1]: Fri Mar 25 13:12:56 2011
Slot 7 (CP1, active): Firmware is being downloaded to the switch. This step may take
up to 30 minutes.

[2]: Fri Mar 25 13:33:23 2011
Slot 7 (CP1, active): Firmware has been downloaded to the secondary partition of the
switch.

[3]: Fri Mar 25 16:38:55 2011
Slot 7 (CP1, active): Firmwaredownloadstatus is entered. System will reboot and a firmware
commit operation will start upon boot up.

[4]: Fri Mar 25 16:49:12 2011
Slot 7 (CP1, standby): The firmware commit operation has started. This may take up to
```


10 minutes.

[5]: Fri Mar 25 16:58:44 2011

Slot 7 (CP1, standby): Firmware commit operation has started to restore the secondary partition.

[6]: Fri Mar 25 16:58:45 2011

Slot 7 (CP1, standby): The firmware commit operation has started. This may take up to 10 minutes.

[7]: Fri Mar 25 17:05:34 2011

Slot 7 (CP1, standby): The commit operation has completed successfully.

SEE ALSO [firmwareCommit](#), [firmwareDownload](#), [firmwareRestore](#), [firmwareShow](#)

firmwareKeyShow

Displays the public key used for signed firmware validation.

Synopsis `firmwarekeyshow`

Description This command displays the contents of the public key used for validating the integrity of firmware images when signed firmware validation is enabled.

Notes A firmware key should be installed on every switch as a part of the Fabric OS installation. The presence of a firmware key does not imply that the firmware signature is checked during **firmwareDownload**. Signed Firmware Download must be enabled before the public key can be used for signature validation.

Use the **configure** command to enable Signed Firmware Download.

If Signed Firmware Download is enabled, and if the validation succeeds, firmware download proceeds normally. If the firmware is not signed or if the signature validation fails, signed firmware download fails.

Refer to the *Fabric OS Administrator's Guide* for complete details on upgrading or downgrading firmware.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the public key used for signed firmware validation:

```
switch:admin> firmwarekeyshow
-----BEGIN PUBLIC KEY-----
MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDjuQpMk4FrceFvVZ12iAakFNv9
k4ZGhFDMPGIHIems1Ywqdl55U7LTDIVwoViTLevtIDn012r1X1NQ+DORAzvJfkWd
XegkeTn/8wDgHBwotPz4WTd9UGJ9M0Vs52ro1TiukIpsh084LXKgxt+IgdseRCzY
8p8rQZWLpyputx6rgwIDAQAB
-----END PUBLIC KEY-----
```

SEE ALSO [firmwareDownload](#), [configureChassis](#)

firmwareRestore

Restores the former active firmware image.

Synopsis `firmwarerestore`

Description Use this command to restore the former active Fabric OS firmware image. This command can only be run if autocommit was disabled during the `firmwareDownload` process.

After a `firmwareDownload` and a `reboot` (with autocommit disabled), the downloaded firmware becomes active. If you do not want to commit the firmware and want to restore the former firmware, issue the `firmwareRestore` command. After running `firmwareRestore`, you can run `firmwareDownload` again.

This command reboots the system and makes the former firmware active. After the switch reboots, both primary and secondary partitions restore the previous firmware.

This command only takes action if the system is booted after a `firmwareDownload`; otherwise, the process terminates with an error.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To restore the former active firmware image:

```
switch:admin> firmwarerestore
Restore old image to be active ...
Restore both primary and secondary image after reboot.
The system is going down for reboot NOW !!
Broadcast message from root (ttyS0) Fri Oct 22 23:48:54 2010...

Doing firmwarecommit now.
Please wait ...
```

SEE ALSO [firmwareCommit](#), [firmwareDownload](#)

firmwareShow

Displays the firmware version and download history.

Synopsis `firmwareshow`

`firmwareshow --history`

`firmwareshow --help`

Description Use this command to display the firmware versions and the firmware download history. The command shows the firmware versions on both the primary and secondary partitions of the storage device. When this command is issued while a firmware download is in process, an appropriate warning message is displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--history

Displays the contents of the firmware download log. For each firmware download record, the command displays the date and time, the switch name, the slot number, port ID, and firmware version.

--help

Displays the command usage.

Examples To display the firmware versions on a Brocade DCX 8510-8 while a firmware download is in progress:

```
switch:admin> firmwareshow
Slot Name  Appl Primary/Secondary Versions  Status
-----
   6  CP0   FOS      v7.2.0          STANDBY
                v7.2.0
   7  CP1   FOS      v7.2.0          ACTIVE *
                v7.2.0
```

WARNING: Firmwaredownload is in progress.

To display the firmware version on a standalone switch:

```
switch:admin> firmwareshow
Appl      Primary/Secondary Versions
-----
FOS       v7.2.0
          v7.2.0
```

To display the firmware download history:

```
switch:admin> firmwreshow --history  
Firmware version history
```

| Sno | Date & Time | Switch Name | Slot | PID | FOS Version |
|-----|--------------------------|-------------|------|------|-----------------------|
| 1 | Fri Feb 18 12:58:06 2011 | CDCX16 | 7 | 1556 | Fabos Version v7.0.0d |
| 2 | Wed Feb 16 07:27:38 2011 | CDCX16 | 7 | 1560 | Fabos Version v7.0.0c |

SEE ALSO [firmwareDownload](#), [firmwareDownloadStatus](#)

firmwareSync

Synchronizes the firmware from the active control processor (CP) to the standby CP.

Synopsis `firmwaresync [-force]`

`firmwaresync --help`

Description Use this command to manually synchronize the firmware from the active CP to the standby CP. Execute the command without arguments to synchronize the active CP firmware to the standby CP.

Notes This command is applicable only for the dual CP systems. Execution of this command updates the firmware version of the standby CP with that of the active CP. The lowest firmware version on the standby CP can be Fabric OS v6.4.0.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

[-force]

Executes the command without user confirmation. This operand is optional.

--help

Displays the command usage.

Examples To synchronize the standby CP with active CP firmware version:

```
switch:admin> firmwaresync
```

```
This command will copy the firmware on the active CP blade to the
standby CP blade but will require that existing telnet, secure telnet or
SSH sessions to the standby CP blade to be restarted.
```

```
This command may take up to 10 minutes.
```

```
Do you want to continue (Y/N) [Y]:
```

```
Firmwaresync has started.
```

```
.....Firmwaresync has been completed successfully.
```

SEE ALSO [Fabric OS Commands](#)

flow

Creates, manages, and displays flows in Flow Vision.

Synopsis

```

flow --create flow_name -feature feature_list
  port_options frame_options [config_options]

flow --activate flow_name -feature feature_list

flow --deactivate flow_name -feature feature_list

flow --control [flow_name] [-feature feature_list]
  [-deviceIdMode mode | -portIdMode mode]
  [-simport port_num -enable | -disable]
  [-size payload_size | -pattern pattern_string]
  [-enable_wrap | -disable_wrap]

flow --modify flow_name port_options

flow --reset flow_name -feature feature_list

flow --show [flow_name] [-feature feature_list | -allzoned]
  [port_options frame_options]
  [-count iterations | -time interval | -verbose [-domain domain_id]]
  [-sortBy field] [-ctrlcfg]
  [-increase [count] | -decrease [count]]

flow --delete flow_name | all [-force]

flow --help

```

Description

Use this command to perform the following functions:

- Create a flow
- Activate or deactivate features for a flow
- Clear the data collected for a particular feature of a flow
- Change the control parameters for a feature or a flow
- Modify the definition of the predefined Analytics VTAP flow
- Display the flow statistics
- Delete a flow

Each of these functions is documented in a separate section that includes function, synopsis, description, operands, and examples.

A flow is a set of related Fibre Channel (FC) frames or packets that share similar traits, such as an ingress port, egress port, or frame options that can uniquely differentiate one set of related frames or packets from a different set of frames. A flow is defined by a combination of ingress port, egress port, source device, and destination device parameters. Flow Vision supports Flow Monitor, Flow Generator, and Flow Mirror features to monitor, simulate, and capture the network traffic respectively. You can create up to 512 flows on Brocade director and 128 flow on switch-based system. For more information on Flow Vision, refer to the *Flow Vision Administrator's Guide*.

Notes This command requires a Fabric Vision license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Function Creating a flow

Synopsis `flow --create flow_name -feature feature_list
port_options frame_options [config_options]`

Description Use this command to create a flow and by default activate the flow for the specified features. This command validates the flow definition before creating a flow. If the flow definition is not supported by any specified feature or if the flow definition is a duplicate of an existing active flow, the operation will terminate with an appropriate error message. For more information on duplicate flows, refer to the *Fabric OS Administrator's Guide*.

The following rules apply for creating a flow:

- You must specify at least one port option (**-ingrport** or **-egrport**) but not both.
- You must specify the source device (**-srcdev**), the destination device (**-dstdev**), or both. The source device and the destination device can be learned using a wildcard ("*").
- The ingress port and egress port must reside in the local switch.
- The source device and destination device can reside in either the local or remote switch.
- You can create 512 flows on a Brocade director and 128 flows on a switch-based system. Note that each flow discovered when using "*" as a source device or destination device will count as one of these flows. These system-discovered flows will impact the number of user-defined flows allowed to be active.
- The configuration options **-noactivate**, **-noconfig**, and **-bidir** are optional.
- You can create up to 64 monitor flows and four generator flows per port.
- You can activate one mirror flow per chassis or switch-based system.

Operands The **--create** command has the following operands:

flow_name

Specifies the name of the flow. Each flow name must be a unique string composed of a maximum of 20 alphanumeric or underscore characters. The flow name is case-insensitive and is always stored as lowercase.

-feature feature_list

Specifies a comma-separated list of features to activate for the flow. Specify "all" to activate all features for a flow. Valid values for *feature_list* include the following:

monitor

Activates the Flow Monitor feature for a flow. The Flow Monitor feature monitors and provides statistics for the specified flow.

The following configuration considerations apply for the Flow Monitor feature:

- The ingress port and egress port can be an SIM Port, F_Port, E_Port, or EX_Port (Gen 5 or later) on the local switch. The ingress port must be a Gen 5 or later port and egress port can be a Gen 4, Gen 5 or later port. The SIM Port must be a Gen 5 or later port.
- VE_Port, VEX_Port, FCoE ports, and GigE are not supported.
- The source device and destination device can be a host or target switch, or a wildcard (""). Using the wildcard ("") enables Flow Vision to learn (discover) all flows on an F_Port and displays statistics for each learnt flow. The learning mode is supported only on Gen 5 or later F_Ports and trunked F_Ports.
- You can use WWN or Fibre Channel ID (FCID) for the **-srcdev** and **-dstdev** operands when creating Flow Monitor on EX_Ports. If the Flow Monitor feature is configured on the ingress port using FCID addressing mode, you must specify the real SID of the source device and the proxy ID of the destination device. If the Flow Monitor feature is configured on the egress port using FCID addressing mode, you must specify the proxy ID of the source device and real DID of the destination device.
- Inter Fabric Link (IFL) flows can be monitored only on EX_Ports in a Fibre Channel router (FCR).
- You can activate the Flow Monitor feature for a maximum of 64 flows per port.
- Access Gateway (AG) mode is supported

generator

Activates the Flow Generator feature for a flow. The Flow Generator is a traffic flow diagnostics feature that generates traffic at line rate in the fabric to validate connectivity, hardware components, and network performance. You must configure the ingress or egress ports and the source and destination devices to be SIM ports before activating a flow using the Flow Generator feature.

You can modify the frame size and payload pattern using the **--control** options after creating the flow.

The following configuration considerations apply for the Flow Monitor feature:

- The Flow Generator feature is not supported on a base switch or ICL port.
- AG mode is not supported.
- You can activate the Flow Generator feature on four flows per port.
- The **-frametype**, **-lun**, and **-bidir** options are not supported with the Flow Generator feature.

mirror

Activates the Flow Mirror feature for a flow. The Flow Mirror captures network traffic in a flow and sends it to the switch CPU for traffic analysis. This feature mirrors only the flows that are in active state. The logging mechanism of the frames can be modified by using the **--control -enable_wrap | -disable_wrap** options.

For both switch-based system and a chassis-based system, a maximum of 256 frames per second are mirrored. A maximum of 1280 frames are stored for switch-based systems and 5120 frames are stored for chassis-based systems.

The following configuration considerations apply for the Flow Mirror feature:

- The ingress port or the egress port must be an F_Port, which is 8G or less, on the local switch.
- The Flow Mirror feature is supported only on Gen 5 or later ports.
- EX_Ports, XISL ports, DISL ports, E_Ports, trunked E_Ports, and trunked F_Ports are not supported.
- AG mode is not supported.
- The Flow Mirror feature is not supported on SIM ports.
- The Flow Mirror feature can be active on only one flow per chassis or switch-based system.

port_options

Specifies the options to configure the ingress or egress ports.

-ingrport port

Specifies the ingress port. The *port* can be in either *[slot/]port* or D,I (Domain, Index) format depending on the port mode configured using the **-portIdMode** operand.

-egrport port

Specifies the egress port. The *port* can be in either *[slot/]port* or D,I format depending on the port mode configured using the **-portIdMode** operand.

The **-ingrport** and **-egrport** options are mutually exclusive; that is, you can specify only one of the options. The learning mode ("") is not supported.

-mirrorport port

Specifies the port from which mirrored frames exits. The *port* can be in either *[slot/]port* or D,I format depending on the port mode configured using the **-portIdMode** operand.

frame_options

Specifies the options that represent the content of the FC frame or attributes of the frame.

-srcdev device_id

Specifies the 3-byte source ID of the originator device. The *device_id* can be in PID or PWWN format depending on the device ID mode configured using the **-deviceIdMode** operand. The PID format is OxDDAAPP, where DD is the Domain ID, AA is the Area ID and PP is the AL_PA ID. The PWWN format is "XX:XX:XX:XX:XX:XX:XX:XX". Specify "" to enumerate all originator devices.

-dstdev device_id

Specifies the 3-byte destination ID of the destination device. The *device_id* can be in PID or PWWN format depending on the device ID mode configured using the **-deviceIdMode** operand. The PID format is OxDDAAPP, where DD is the Domain ID, AA is the Area ID and PP is the AL_PA ID. The PWWN format is "XX:XX:XX:XX:XX:XX:XX:XX". Specify "" to enumerate all destination devices.

-lun LUN_id

Specifies the LUN identification in decimal or hexadecimal (2-byte value) format. The valid range for *LUN_id* is from 0 through 65535 (0x0000 to 0xFFFF). The LUN parameter includes the addressing method and fields associated with the respective addressing method. If the most significant byte (MSB) is not specified in the flow definition, 0x00 is used. This operand is not supported by the Flow Generator feature.

-frametype type

Specifies the predefined frame type. Valid values for *type* include the following: SCSI, SCSIRead, SCSIWrite, SCsIRW, SCSI2Reserve, SCSI3Reserve, scsicmdsts, ABTS, BAACC, SCSI2Release, SCSI3Release, SCsITur, SCSI2ReserveRelease, SCSI3ReserveRelease, SCsIGoodStatus, SCsICheckStatus, SCsIResvConflict, SCsIIquiry, SCsIXferdy, Srr, and BARJT. This operand is not supported by the Flow Generator feature.

config_options

Specifies the flow configuration options.

-bidir

Configures the features to act on both incoming and outgoing directions on bidirectional traffic. For example, the Flow Mirror feature mirrors the frames on both directions: traffic originating from the source device and traffic that is destined to the source device. This operand is not supported by the Flow Generator feature.

-noactivate

Allows a flow to be created, but suppresses activation of a flow.

-noconfig

Overrides the default behavior, which automatically saves the flow in the persistent configuration when you create a flow. The flow will be lost when the system is rebooted.

Examples

To create a flow on ingress port (1/10) for traffic from the source device (0x020a00) to the destination device (0x030000) and activate the flow for the Flow Mirror and Flow Monitor features:

```
switch:admin> flow --create myflow1 -feature mirror,monitor
                -ingrport 1/10 -srcdev 0x020a00 -dstdev 0x030000
```

To create a flow to monitor SCSI frames going through an egress port:

```
switch:admin> flow --create scsicsflow -feature monitor
                -egrport 1/9 -frametype scsichkstatus
```

To create a flow to monitor statistics from a particular source to destination on a specified LUN:

```
switch:admin> flow --create lunflow -feature monitor
                -srcdev 0x010502 -dstdev 0x030700 -ingrport 1/5 -lun 4
```

To create a flow from source device (0x010203) to destination device(0x020304) for Flow Generator feature without saving the flow:

```
switch:admin> flow --create myflow2 -feature generator
                -srcdev 0x010203 -dstdev 0x020304 -ingrport 2/3 -noconfig
```

To create a flow from source device (0x01ba00) to destination device(0x01c000) for Flow Generator feature without activating the flow:

```
switch:admin> flow --create myflow3 -feature generator
                -srcdev 0x01ba00 -dstdev 0x01c000 -ingrport 2/18 -noactivate
```

To create a flow to monitor bidirectional traffic from a particular source to destination:

```
switch:admin> flow --create endtoendflow -feature monitor
                -ingrport 2/5 -srcdev 010500 -dstdev 040900 -bidir
```

To create a flow to monitor traffic from source devices to destination devices in learning mode:

```
switch:admin> flow --create flowLearn -feature monitor
                -ingrport 3/5 -srcdev "" -dstdev ""
```

To create a flow to mirror ingress and egress traffic flowing through a port to an embedded port:

```
switch:admin> flow --create myflow4 -feature mirror
                -ingrport 1/20 -srcdev 0x010203 -lun 0x02 -bidir
```

To create a flow to generate traffic from the source device (0x020a00) to all destination devices that are zoned with the source device:

```
switch:admin> flow --create myflow5 -feature generator
                -ingrport 1/10 -srcdev 0x020a00 -dstdev ""
```

Function Activating features or flow

Synopsis `flow --activate flow_name -feature feature_list`

Description Use this command to activate a flow if it is not in the active state or configure the specified features for a flow. Use this command if you created a flow using the **-noactivate** operand or if none of the features are activated for the flow. Note that the flows are not enforced until they are active and the flows cannot be activated without configuring the features. If one or more of the flow parameters are not supported by any specified feature, the flow will not be activated and the operation will fail with an appropriate error message.

Operands The **--activate** command has the following operands:

flow_name

Specifies the name of the flow on which the specified features must be activated. Specify "all" to activate the specified features for all flows. The following predefined flows are supported:

sys_mon_all_fports

Used for automatically learning and monitoring the traffic passing through all the online F_Ports in the switch. By default, this flow is in the deactivated state. This flow is supported only with the **monitor** feature.

sys_analytics_vtap

Used to mirror frames to a port on a remote domain, which is in IO Analyzer Mode. This flow is supported only with the **mirror** feature.

feature_list

Specifies a comma-separated list of features to activate for a flow. Specify "all" to activate all features for a flow. Valid values for *feature_list* include the following: **monitor**, **generator**, and **mirror**.

Examples To activate Flow Mirror feature for a flow:

```
switch:admin> flow --activate myflow2 -feature mirror
```

Function Deactivating features

Synopsis `flow --deactivate flow_name -feature feature_list`

Description Use this command to deactivate the specified features for a flow. All data associated with the deactivated feature will be cleared.

Operands The `--deactivate` command has the following operands:

flow_name

Specifies the name of the flow on which the specified features must be deactivated. Specify "all" to deactivate the specified features for all flows. The following predefined flow is supported:

sys_mon_all_fports

Used for automatically learning and monitoring the traffic passing through all the online F_Ports in the switch. This flow is supported only with the **monitor** feature.

feature_list

Specifies a comma-separated list of features to deactivate for the flow. Specify "all" to deactivate all features for a flow. Valid values for *feature_list* include the following: **monitor**, **generator**, and **mirror**.

Examples To deactivate the Flow Mirror feature for a flow:

```
switch:admin> flow --deactivate myflow2 -feature mirror
```

Function Changing the control parameters

Synopsis `flow --control [flow_name] [-feature feature_list]
[-deviceIdMode mode | -portIdMode mode]
[-simport port_num -enable | -disable]
[-size payload_size | -pattern pattern_string]
[-enable_wrap | -disable_wrap]`

Description Use this command to modify the port or device addressing mode and to change the control parameters of a flow or features.

Operands The `--control` command has the following operands:

flow_name

Specifies the flow name. Specify "all" to modify the control parameters for all flows. If a flow name is not specified, the default values of the control parameters, `-size` and `-pattern` will be modified; the flows which are already present does not intake the modified control parameters. The modified control parameters will apply to all flows created in future.

`-feature feature_list`

Specifies a comma-separated list of features on which the control parameters should be applied. Valid values for *feature_list* include the following: **generator** and **mirror**.

`-deviceIdMode mode`

Specifies the mode for addressing the devices. The valid values for *mode* are **pid** and **wwn**. The default mode is **pid**.

`-portIdMode mode`

Specifies the mode for addressing the ports. The valid values for *mode* are **slotport** and **index**. The default mode is **slotport**.

`-simport port`

Enables or disables a port as a SIM port. The *port* can be in `[slot/]port` or D,I format depending on the port ID mode configured using the `-portIdMode` operand. You can specify a single port or a range of ports in `slot/port`, `slot/port-port`, `slot/*` (all ports in slot), `*/*` (all ports in switch or logical switch), `port`, `port-port`, and `*` (all ports in switch or logical switch) formats. The flow name is not required for enabling or disabling a port as SIM port.

`-enable`

Enables the port as a SIM port.

`-disable`

Disables the SIM port configuration on the port.

`-size payload_size`

Specifies the size of the frame payload. The frame size must be specified in multiples of 4, for example, 64, 68, 120, or 320 characters. The valid range is from 64 through 2048 characters; 0 for random size. The default payload size is 2048 bytes.

The following operands are specific to the flow **generator** feature:

`-pattern pattern_string`

Specifies the ASCII pattern of the payload. The valid range is from 1 through 32 bytes; 0 for random pattern. The default is random pattern.

The following operands are specific to the flow **mirror** feature. You can use these operands only if the Flow Mirror feature is not activated for a flow. The flow name is not required for enabling or disabling this option.

-enable_wrap

Enables the wrapping of mirrored frame logs. This is the default value. The Flow Mirror flow must be deactivated before using this command. The statistics of the Flow Mirror flow are reset during this operation.

-disable_wrap

Disables the wrapping of mirrored frame logs. The Flow Mirror flow must be deactivated before using this command. The statistics of the Flow Mirror flow are reset during this operation.

Examples

To change the device addressing mode to WWN:

```
switch:admin> flow --control -deviceidmode wwn
```

To change the port addressing mode to index:

```
switch:admin> flow --control -portidmode index
```

To configure the ports from 1 to 40 in slot 1 as SIM port:

```
switch:admin> flow --control -simport 1/1-40 -enable
```

To remove the SIM port configuration on port 20 in slot 1:

```
switch:admin> flow --control -simport 1/20 -disable
```

To modify the size and pattern for a flow:

```
switch:admin> flow --control myflow2 -feature generator
-size 100 -pattern "1234"
```

To disable wrapping of frame logs for the Flow Mirror feature:

```
switch:admin> flow --control -feature mirror -disable_wrap
```

Function Modifying definition of Analytics vTap flow

Synopsis `flow --modify flow_name port_options`

Description Use this command to modify the existing flow definition of the predefined flow "sys_analytics_vtap" and specify either a single port, a MAPS logical port group name, or a wildcard (*) as ingress port. You can also override the default mirror port by using the `--mirrorport port` option.

Operands The `--modify` command has the following operands:

flow_name

Specifies the name of the Analytics vTap flow. Valid option is **sys_analytics_vtap**. The `sys_analytics_vtap` flow is used to mirror frames to a port on a remote domain, which is in IO Analyzer Mode. The flow can be modified only if it is in the deactive state.

port_options

Specifies the options to configure the ingress or mirror port.

-ingrport *port*

Specifies the ingress port. The *port* can be in either [*slot/*]*port* or D,I (Domain, Index) format depending on the port mode configured using the **-portIdMode** operand, or a wildcard (*).

-mirrorport *port*

Specifies the port from which mirrored frames exits. The *port* must be in the D,I (Domain, Index) format.

Examples To modify and activate the predefined Analytics vTap flow:

```
switch:admin> flow --modify sys_analytics_vtap -ingrport 2/10 -mirrorport 12, 21
switch:admin> flow -activate sys_analytics_vtap -feature mirror
```

Function Clearing feature data

Synopsis `flow --reset flow_name -feature feature_list`

Description Use this command to clear the data of the specified flow and feature combination. All the data collected for the flow and feature combination will be cleared.

Operands The **--reset** command has the following operands:

flow_name

Specifies the name of the flow on which the feature data must be cleared. Specify "all" to reset all flows. The following predefined flow is supported:

sys_mon_all_fports

Used for automatically learning and monitoring the traffic passing through all the online F_Ports in the switch. This flow is supported only with the **monitor** feature.

feature_list

Specifies a comma-separated list of features. Specify "all" to reset all features for a flow. Valid values for *feature_list* include the following: **monitor**, **generator**, and **mirror**.

Examples To clear all the mirrored frames for a flow:

```
switch:admin> flow --reset myflow1 -feature mirror
```

Function Displaying flow data

Synopsis `flow --show [flow_name] [-feature feature_list | -allzoned]
[port_options frame_options]
[-count iterations | -time interval | -verbose [-domain domain_id]]
[-sortBy field] [-ctrlcfg]
[-increase [count] | -decrease [count]]`

Description Displays the flow definitions and frame statistics. When used without operand, the **--show** command displays the list of all flows configured in Flow Vision.

The show on demand feature allows you to display the statistics of the features without creating a flow. To display the statistics of the features without creating a flow, you must specify the operands without a flow name and provide a flow definition. The flow will run for a time interval of 6 seconds and the data will be displayed for the specified features.

Operands The **--show** command has the following operands:

flow_name

Displays the parameters and flow statistics for the specified flow. Specify "all" to display all flows. The following predefined flows are supported:

sys_mon_all_fports

Used for automatically learning and monitoring the traffic passing through all the online F_Ports in the switch. By default, this flow is in the deactivated state. You can activate this flow using the **--activate** operand. This flow is supported only with the **monitor** feature.

sys_analytics_vtap

Used to mirror frames to a port on a remote domain, which is in IO Analyzer Mode. This flow is supported only with the **mirror** feature.

-feature feature_list

Displays the statistics for the specified features. Specify "all" to display statistics of all configured features. Valid values for *feature_list* include the following: **monitor**, **generator**, **mirror**, and **fabinfo**.

fabinfo

Displays the topology data (the summary of all the paths taken by the flows in the flow definition) and the MAPS violations. The *flow_name*, **-lun**, **-frametype**, and **-sortby** options are not supported with this feature.

-allzoned

Displays all the zoned devices for a specified port, flow name, or flow definition.

port_options

Specifies the ingress or egress port options.

-ingrport port

Specifies the ingress port. The *port* can be in either *[slot/]port* or D.I (Domain, Index) format depending on the port mode configured using the **-portIdMode** operand.

-egrport port

Specifies the egress port. The *port* can be in either *[slot/]port* or D.I format depending on the port mode configured using the **-portIdMode** operand.

-mirrorport port

Specifies the port from which mirrored frames exits. The *port* can be in either *[slot/]port* or D.I format depending on the port mode configured using the **-portIdMode** operand.

frame_options

Specifies the options that represent the content of the FC frame or attributes of the frame.

- srcdev *device_id***
Specifies the 3-byte source ID of the originator device. The *device_id* can be in PID or PWWN format depending on the device ID mode configured using the **-deviceIdMode** operand. The PID format is OxDDAAPP, where DD is the Domain ID, AA is the Area ID and PP is the AL_PA ID. The PWWN format is "XX:XX:XX:XX:XX:XX:XX:XX". Specify "" to enumerate all originator devices.
- dstdev *device_id***
Specifies the 3-byte destination ID of the destination device. The *device_id* can be in PID or PWWN format depending on the device ID mode configured using the **-deviceIdMode** operand. The PID format is OxDDAAPP, where DD is the Domain ID, AA is the Area ID and PP is the AL_PA ID. The PWWN format is "XX:XX:XX:XX:XX:XX:XX:XX". Specify "" to enumerate all destination devices.
- lun *LUN_id***
Specifies the LUN identification in decimal or hexadecimal (2-byte value) format. The valid range for *LUN_id* is from 0 through 65535 (0x0000 to 0xFFFF). The LUN ID includes the addressing mode. If the most significant byte (MSB) is not specified, 0x00 is used. This operand is not supported by the Flow Generator feature.
- frametype *type***
Specifies the predefined frame type. Valid values for *type* include the following: SCSI, SCSIRead, SCSIWrite, SCISIRW, SCSI2Reserve, SCSI3Reserve, scsicmdsts, ABTS, BAACC, SCSI2Release, SCSI3Release, SCISITur, SCSI2ReserveRelease, SCSI3ReserveRelease, SCISIGoodStatus, SCISICheckStatus, SCISIRsvConflict, SCISIIquiry, SCISIXferdy, Srr, and BARJT. This operand is not supported by the Flow Generator feature.
- count *value***
Specifies the number of times the output must be repeated. The default value is 1 and the maximum value is 10.
- time *time_interval***
Displays the flow information for the specified time interval until it is terminated with **Ctrl + C**. Values are in seconds. The time interval must be specified in multiples of 5, for example, 5, 10, 15, or 20 seconds.
- verbose**
Displays the flow or feature statistics in verbose mode.
- domain *domain_id***
Displays detailed information for the specified domain. The operand is valid only with the **fabInfo** feature.
- sortby *field***
Displays the sub-flows for a feature in the specified order. The *field* can be columnX, where X is the column number. The value of X can be from 1 to maximum number of columns present in the **--show** output. For the Flow Mirror feature, the sub-flows may or may not be present but the mirrored frames can be sorted. The **-sortby** option can be applied if only one feature is specified with the **--show** command.
- ctrlcfg**
Displays the control parameters.

-increase [count]

Displays the monitoring statistics in ascending order based on the throughput. The *count* parameter is optional and it can be used to specify the number of entries to be displayed. This operand is supported only with the **monitor** feature.

-decrease [count]

Displays the monitoring statistics in descending order based on the throughput. The *count* parameter is optional and it can be used to specify the number of entries to be displayed. This operand is supported only with the **monitor** feature.

Examples To display all flows in Flow Vision:

```
switch:admin> flow --show
----- \
  Flow Name | Feature | SrcDev | DstDev | IngrPt | EgrPt | \
----- \
myflow1    | mon+   | 020a00 | 030000 | 1/10   | -     | \
scsicsflow | mon+   | 520a00 | 520b00 | -      | 1/9   | \
lunflow    | mon+   | 010502 | 030700 | 1/5    | -     | \
myflow2    | gen+   | 010203 | 020304 | 2/3    | -     | \
myflow3    | gen    | 01ba00 | 01c000 | 1/18   | -     | \
endtoendflow | mon+ | 010500 | 040900 | 2/5    | -     | \
flowLearn  | mon+   | *      | *      | 3/5    | -     | \
myflow4    | mir+   | 010203 | -      | 1/20   | -     | \
myflow5    | gen+   | 020a00 | *      | 1/10   | -     | \
----- \
  BiDir | LUN | FrameType |
----- \
  no    | -   | -         |
  no    | -   | sstatc    |
  no    | 4   | -         |
  no    | -   | -         |
  no    | -   | -         |
  yes   | -   | -         |
  no    | -   | -         |
  no    | 0x02 | -         |
  no    | -   | -         |
+ Denotes feature is currently activated for the flow
```

To display the flow definition and statistics of a flow:

```
switch:admin> flow --show myflow2
=====
Name      : myflow2      Features: gen(Activated)
Definition: IngrPort(3),SrcDev(0x010203),DstDev(0x020304)

Flow Generator (Activated):
-----
| SrcDev | DstDev |
-----
| 0x010203 | 0x020304 |
-----
Number of frames generated from IngrPort : 1.50G
=====
```

To display the statistics of the Flow Monitor feature:

```
switch:admin> flow --show myflow1 -feature monitor
=====
Name      : myflow1      Features: mon(Activated)
Definition: IngrPort(10),SrcDev(0x020a00),DstDev(0x030000)

Flow Monitor (Activated):
Monitor time: | Tue Jun 25 04:20:06 PDT 2013 |
-----
|      Frame Count      |      Frames Per Sec.      |      Byte count      | \
| Tx / Rx / Total | Tx / Rx / Total | Tx / Rx / Total | \
| 18.20G/ 0 / 18.20G| 786.52k/ 0 / 786.52k| 73.48T/ 0 / 73.48T | \
-----
===== \

|      Throughput(Bps)      |      Frame Size(Bytes) |
| Tx / Rx / Total | Tx / Rx |
| 1.56G/ 0 / 1.56G| -- / -- |
-----
=====
```

To display the monitor statistics from a particular source to destination on a specified LUN:

```
switch:admin> flow --show lunflow
=====
Name : lunflow Features: mon(Activated)
Definition: IngrPort(5),SrcDev(0x010502),DstDev(0x030700),Lun(4)
Flow Monitor:
Monitor time: | Thu Jun 06 15:15:39 UTC 2013 |
-----
| I/O Count      | I/O Per Sec.(IOPS)      | I/O bytes Transferred | \
| Reads / Writes/ Total | Reads / Writes/ Total | Reads / Writes/ Total | \
| 44.92k/ 44.94k/ 89.85k| 2.00k/ 2.00k/ 4.01k | 5.88M/ 5.89M/ 11.77M | \
-----
===== \

| I/O bytes Per Sec. |
| Reads / Writes/ Total |
| 2.62M/ 2.62M/ 5.25M |
-----
=====
```

To display the Flow Monitor statistics for bidirectional traffic:

```
switch:admin> flow --show endtoendflow -feature monitor
=====
Name : endtoendflow Features: mon(Activated)
Definition: IngrPort(1,5),SrcDev(010500),DstDev(040900),BiDir
Flow Monitor:
```

```

Monitor time: | Thu Jun 06 15:15:39 UTC 2013 |
-----
| Frame Count          | Frames Per Sec.      | Byte count           |
| Tx / Rx / Total     | Tx / Rx / Total     | Tx / Rx / Total     |
-----
| 2.92M/ 2.96M/ 5.88M | 159.31k/161.78k/321.10k | 5.99G/ 5.99G/ 11.99G |
-----
| Throughput(Bps)      | Frame Size(Bytes)   |
| Tx / Rx / Total     | Tx / Rx             |
-----
| 327.15M/327.28M/654.44M | 2052 / 2020         |
-----
| I/O Count           | I/O Per Sec.(IOPS) | I/O bytes Transferred |
| Reads / Writes/ Total | Reads / Writes/ Total | Reads / Writes/ Total |
-----
| 44.92k/ 44.94k/ 89.85k | 2.00k/ 2.00k/ 4.01k | 5.88G/ 5.89G/ 11.77G |
-----
=====
| I/O bytes Per Sec.  |
| Reads / Writes/ Total |
| 262.84M/262.73M/525.68M |
-----
=====

```

To display the monitor statistics for a flow:

```

switch:admin> flow --show flowLearn -feature monitor
-----
| SID(*) | DID(*) | Rx Frames Count | Rx Frames per Sec. | Rx Bytes Count |
-----
| 010500 | 030700 | 2.92M           | 159.31k            | 5.99G           |
| 010501 | 030700 | 1.92M           | 100.22k            | 4.91G           |
| 010502 | 030700 | 0.92M           | 51.81k             | 2.36G           |
| 010500 | 040900 | 2.92M           | 259.31k            | 7.89G           |
| 010501 | 040900 | 1.92M           | 159.31k            | 3.72G           |
| 010502 | 040900 | 0.92M           | 159.31k            | 1.08G           |
-----
| *      | *      | 12.72M         | 770.37K            | 25.86G          |
-----
| Rx Throughput(Bps) | Avg Rx Frm Sz(Bytes) |
-----
| 627.15M            | 2052                  |
| 444.62M            | 1052                  |
| 314.17M            | 452                   |
| 327.15M            | 2052                  |
| 327.15M            | 1052                  |
| 327.15M            | 452                   |
-----
| 2.68G              | 1192                  |
-----

```

To display the statistics for the Flow Mirror feature in learning mode:

```
switch:admin> flow --show mirFlow -feature mirror
Name : mirFlow      Features: mir(Activated)
Definition: EgrPort(17),DstDev(0x051100),BiDir
Flow Mirror:
----- \
| SID(*) | DID(*) | OXID | RXID | SOF  | EOF  | Frame_type \
----- \
| 051600 | 051100 | 0883 | ffff | SOFi3 | EOFt | SCSIRead   \
| 051600 | 051100 | 0191 | 0bce | SOFn3 | EOFn | Data       \
| 051600 | 051100 | 0c8f | ffff | SOFi3 | EOFt | SCSIWrite  \
| 051600 | 051100 | 0017 | ffff | SOFi3 | EOFt | SCSIWrite  \
| 051600 | 051100 | 0191 | 0bce | SOFn3 | EOFn | Data       \
----- \
| Dir | Time-Stamp |
----- \
| Tx | Jun 10 11:08:10 |
| Tx | Jun 10 11:08:10 |
| Tx | Jun 10 11:08:10 |
| Tx | Jun 10 11:08:10 |
| Tx | Jun 10 11:08:10 |
----- \
```

To display all the devices zoned with device at port '11' (F_Port) as destination:

```
switch:admin> flow --show -allzoned -src "" -dst Ox190b00 -egr 11
Active Flows:
----- \
| SrcDev | DstDev |
----- \
| 0x21400 0x21e00 0x21f00 0x31000 0x31100 | 0x190b00 |
----- \
```

To display all the devices zoned with device at port 11 (F_Port) as source:

```
switch:admin> flow --show -allzoned -src "" -dst * -ing 11
Active Flows:
----- \
| SrcDev | DstDev |
----- \
| 0x190b00 | 0x21400 0x21e00 0x21f00 0x31000 0x31100 |
----- \
```

To display all the zoned devices for port 15 (E_Port):

```
switch:admin> flow --show -allzoned -src "" -dst "" -ing 15
Active Flows:
----- \
| SrcDev | DstDev |
----- \
| 0x190a00 | 0x21400 0x21e00 0x21f00 0x31000 0x31100 |
----- \
| 0x190b00 | 0x21400 0x21e00 0x21f00 0x31000 0x31100 |
----- \
```

To display all the zoned devices for a flow:

```
switch:admin> flow --create flow3 -fe mon -src "" -dst Ox190b00 -egr 11
switch:admin> flow --show flow3 -allzoned
Active Flows:
-----
| SrcDev                               | DstDev |
-----
| 0x21400 0x21e00 0x21f00 0x31000 0x31100 | 0x190b00 |
-----
```

To display the statistics for the Flow Mirror feature in verbose mode:

```
switch:admin> flow --show mirFlow1 -feature mirror -verbose
Name : mirFlow1 Features: mir(Activated)
Definition: EgrPort(1/5),SrcDev(010200),DstDev(*),BiDir
Flow Mirror:
-----
Time-Stamp      | Dir | SOF | EOF | Frame_Type | Frame Contents |
-----
Jun 04 08:27:04 | Tx  | SOFi3 | EOFt | SCSI3_Res | 06040500 00010200 08290000
009a0fd3 00000000 00000000 00000000 00000001 5f010300 00000000 18000000 00000000
Jun 04 08:27:04 | Tx  | SOFi3 | EOFt | SCsITxRdy | 05040500 00010200 08890000
07ce01ca 00000000 00000000 00000018 00000000
Jun 04 08:27:04 | Rx  | SOFi3 | EOFt | Data | 01010200 00040500 08090008
089a0f71 00000000 00000000 ffaabbcc 00000000 00000000 00000000 00000000
Jun 04 08:27:04 | Rx  | SOFi3 | EOFt | SCsIGoodSts | 07010200 00040500 08990000
089a0f71 00000000 00000000 00000000 00000000 00000000 00000000 00000000
```

To display the feature statistics in sorted order:

```
switch:admin> flow --show myflow -feature monitor -sortBy column1
Monitor time: | Fri Jun 07 03:56:24 MDT 2013 |
=====
Name          : myflow Features: gen(Activated), mon(Activated)
Definition: EgrPort(4,8),SrcDev(0x*),DstDev(0x040800)
Flow Monitor:
----- \
|SID(*)| Tx Frames Count | Tx Frames per Sec. | Tx Bytes Count | \
----- \
| 40600|      11.99M     |      169.68k       |      25.08G     | \
| 40700|      11.99M     |      169.86k       |      25.08G     | \
| 40900|      28.49M     |      403.22k       |      59.38G     | \
----- \
| Tx Throughput(Bps) | Avg Tx Frm Sz(Bytes)|
----- |
|      354.97M       |      2088           |
|      355.35M       |      2088           |
|      124.48M       |      2080           |
----- |
```

To display the SCSI frame statistics for the Flow Monitor feature:

```
switch:admin> flow --show scsicsflow -feature monitor
-----
Frame Statistics
-----
| Tx Frames Count | Tx Frames per Sec. |
-----
| 238             | 5                   |
-----
```

To display the control parameters:

```
switch:admin> flow --show -ctrlcfg
SimPort Information
-----|-----|-----|-----|-----|-----|
Slot | Port | PID | PWWN | SID Frame Count | \
-----|-----|-----|-----|-----|-----|
1    | 2    | 050200 | 20:02:00:05:1e:e2:8e:00 | 0K | \
      |-----|
      | DID Frame Count |
      |-----|
      | 19.46K |
      |-----|
```

To display the control parameters for the Flow Generator feature:

```
switch:admin> flow --show -ctrlcfg -feature generator
Control Parameters of Generator
Size: 1024
Pattern: SCSITur
```

To display the statistics for the Flow Mirror feature on ingress port (1/10) without creating a flow (show on demand):

```
switch:admin> flow --show -srcdev 0x030000 -ingrport 1/10 -feature mirror
```

To display the flow dashboard data:

```
switch:admin> flow --show -feature fabinfo -srcdev "" -egrport 17
Flow Dashboard Information:
=====

Topology Data:
-----

srcDev (0x170500), dstDev(0x31100)
srcDev(0x170500) <-> (23/1#, 23/5) <<-> (3/7, 3/17) <-> dstDev(0x31100)

srcDev (0x190a00), dstDev(0x31100)
[srcDev: 0x190a00]<-> (25/10, 25/2)<->(23/2#, 23/5)<->(3/7, 3/17#)<->
 [dstDev:0x31100]
[srcDev:0x190a00]<-> (25/10, 25/3)<->(21/3, 21/1#)<->(3/4, 3/17#)<-> [dstDev:0x31100]
[srcDev:0x190a00]<-> (25/10, 25/4)<->(21/4, 21/1#)->(3/4, 3/17#)<-> [dstDev:0x31100]

# "Indicates there are MAPS violations on these ports"

Switch Specific Data:
=====
```



```
SwitchDomain      : 3 (0x3)
Name: sw0  Model  : 121.3  Uptime: (41 days 0 hrs 02 mins)
FirmwareVersion  : v8.0.0
OperationalStatus : Healthy
RebootReason     : Reboot
Fenced Ports    : 24
Decommissioned Ports : None
Quarantined Ports : None
```

MAPS violation:

```
-----
Port Health(24) |6 | defALL_HOST_PORTSCRC_2|09/17/14 09:13:24|Port17| 5 |
Port Health(15) |2 | defNON_E_F_PORTSFLF_0 |09/17/14 09:18:24|Port17| 3 |
```

```
=====
SwitchDomain: 25 (0x19)
Name: sw0  Model: 121.3  Uptime: (40 days 21 hrs 12 mins)
FirmwareVersion      : v8.0.0
OperationalStatus    : CRITICAL Reason : FAULTY_BLADE
RebootReason         : Reboot
Fenced Ports         : None
Decommissioned Ports : None
Quarantined Ports    : None
```

```
Category(Rule Count)|RepeatCount|Rule Name |Execution Time |Object |Triggered
Value(Units)|
```

No violations on the ports in the flow path

```
=====
SwitchDomain: 21 (0x15)
Name: sw0  Model: 121.3  Uptime: (40 days 03 hrs 10 mins)
FirmwareVersion      : v8.0.0
OperationalStatus    : Healthy
RebootReason         : Reboot
Fenced Ports         : None
Decommissioned Ports : None
Quarantined Ports    : None
```

No MAPS violations

To display the bottom five monitoring entries based on throughput:

```
switch:admin> flow --show sys_mon_all_fports -fe mon -increase 5
```

```
=====
Name      : sys_mon_all_fports  Features: mon(Activated)          noConfig: Off
Definition: IngrPort(*),SrcDev(*),DstDev(*)
```

```
Flow Monitor (Activated):
Monitor time: | Wed Dec 10 23:49:48 UTC 2014 |
```

flow

| Ingr(*) | SID(*) | DID(*) | Rx Frames Count | Rx Frames per Sec. | Rx Bytes Count | Rx Throughput(Bps) | Avg Rx Frm Sz(Bytes) |
|---------|--------|--------|-----------------|--------------------|----------------|--------------------|----------------------|
| 13 | 010d00 | 010a00 | 2.26M | 377.30k | 1.51G | | |
| 259.38M | | | 720 | | | | |
| 14 | 010e00 | 010d00 | 1.98M | 330.81k | 1.74G | | |
| 298.36M | | | 948 | | | | |
| 18 | 011200 | 010c00 | 2.28M | 381.04k | 1.76G | | |
| 300.72M | | | 828 | | | | |
| 20 | 011400 | 011200 | 1.89M | 315.93k | 1.78G | | |
| 303.97M | | | 1008 | | | | |
| 16 | 011000 | 011200 | 1.99M | 333.30k | 1.80G | | |
| 307.82M | | | 968 | | | | |
| * | * | * | 92.41M | 15.40M | 99.69G | | |
| 16.61G | | | 1160 | | | | |

To display flow dashboard information for a specific domain:

```

switch:admin> flow --show -feature fabinfo -srcdev Ox011100 \
                -egrport 17 -verbose -domain 3
Flow Dashboard Information:
=====

Switch Specific Data:

=====

SwitchDomain: 3 (0x3)
Name: sw0      Model: 121.3      Uptime: (41 days 0 hrs 02 mins)
FirmwareVersion : v8.0.0
OperationalStatus: Healthy
RebootReason: Reboot

MAPS violation:
-----
Port Health(24) |6      | defALL_HOST_PORTSCRC_2|09/17/14 09:13:24|Port17| 5 |
Port Health(15) |2      | defNON_E_F_PORTSLF_0 |09/17/14 09:18:24|Port17| 3 |

MAPS history data:
-----
Stats(Units)  Current  05/06/15  05/05/15  05/04/15  04/30/15  04/29/15
04/28/15
Port(val) Port(val) Port(val) Port(val) Port(val) Port(val)
Port(val)
-----
CRC           32(3)   32(2)   46(34)   45(7)    32(1)    -       32(7)
42(2)        42(2)   43(33)   44(6)    33(1)    -       33(7)
12(1)        12(1)   42(31)   42(5)    34(1)    -       34(7)

```

| | | | | | | | |
|-----|--------|-------|----------|---------|-------|------|--------|
| | 06(1) | 06(1) | 47(20) | 46(5) | 35(1) | - | 35(7) |
| ITW | 32(11) | 32(8) | 45(2300) | 45(539) | 32(3) | 0(1) | 32(21) |
| | 33(3) | 33(3) | 40(2282) | 40(522) | 33(1) | - | 33(7) |
| | 34(3) | 34(3) | 44(2276) | 41(519) | 34(1) | - | 34(7) |
| | 35(3) | 35(3) | 41(2269) | 44(517) | 35(1) | - | 35(7) |

Function Deleting a flow

Synopsis `flow --delete flow_name`

Description Use this command to delete the specified flow from Flow Vision. All the data associated with the flow will be cleared.

Operands The `--delete` command has the following operands:

`flow_name` | `all`

Specifies the name of the flow to be deleted. Specify `all` to delete all flows. Use of `all` option prompts for a confirmation.

`-force`

Executes the command without confirmation. This operand is optional.

Examples To delete a flow:

```
switch:admin> flow --delete myflow1
```

To delete all flows:

```
switch:admin> flow --delete all
```

This operation will delete all user defined flows and deactivate pre-defined flows.

Are you sure? (yes, y, no, n): [no] y

All flows deleted successfully.

Function Displaying the command help

Synopsis `flow --help [operand] [...]`

Description Use this command to display the command usage.

Operands `operand`

Displays the syntax and usage guidelines for the specified operand or sub-option.

Examples To display the command usage:

```
switch:admin> flow --help
```

Description - Perform Flow Vision operations

SYNTAX: flow <Flow_operations>

<Flow_operations>

--create <flowname> -feature <list>

<Port options> <Frame options> <Config options>

```

--delete          <flowname>
--control         [flowname] [-feature <list>]
                  [-deviceIdMode <mode> | -portIdMode <mode>]
                  [-simport <portID> -enable/-disable]
                  [-size <frame size> | -pattern <pattern string>]
                  [-enable_wrap/-disable_wrap]
--reset          <flowname> -feature <list>
--activate       <flowname> -feature <list>
--deactivate     <flowname> -feature <list>
--show           [flowname] [-feature <list>]
                  [<Port options> <Frame_options>]
                  [-count <iterations> | -time <interval> | -verbose]
                  [-sortby <field>] [-ctrlcfg]

```

To display the command usage for an operand:

```

switch:admin> flow --help --create
Description - Create a flow definition for Flow Vision features
SYNTAX: flow --create <flowname> -feature <list>
          <Port options> <Frame options> <Config options>

<flowname>          : Unique string of characters to identify a flow
-feature <list>     : Comma separated list of Flow Vision features
                    (mirror, monitor and generator)
<Port options>
  -ingrport <portID>: Switch receive port on which feature(s) is applied
  -egrport <portID>: Switch transmit port on which feature(s) is applied

<Frame options>
  -srcdev <devID>  : FC device that is flow source (SID or PWWN)
  -dstdev <devID>  : FC device that is flow destination (DID or PWWN)
  -lun <LUN ID>    : SCSI LUN identifier
  -frametype <ftyp>: Frame type applicable to the flow

<Config options>
  -bidir           : Applies Flow Vision features for both directions of flow
  -noactivate     : Creates flow without activating the flow
  -noconfig       : Creates flow without saving persistently

```

To display the command usage for an operand and sub-option:

```

switch:admin> flow --help --create -frametype
-frametype <ftyp>: FC frame types applicable to the flow.
                  <ftyp> can be one of scsi, srd, swr, srdwr, sres2, sres3,
                  scsirel2, stur, scsiresrel2, scsigoodstatus,
                  scsichkstatus, scsirel3, scsiresrel3, srescnfl, sinq,
                  sxfr, scmdsts, abts, bacc, barjt.

```

SEE ALSO None

fosConfig

Displays or modifies Fabric OS features.

Synopsis `fosconfig --enable feature`

`fosconfig --disable feature`

`fosconfig --show`

Description Use this command to enable or disable a feature, or to display the current operating status of features on a switch. This command can be run while the switch is online.

The following features are supported (refer to the Notes for limitations):

- FC Routing service (see **fcrConfigure**)
- Virtual Fabrics (see **IfCfg** and **IsCfg**)
- Ethernet switching service

Notes The features described may not be supported on all platforms. If you attempt to enable a feature that is not supported on your platform, an error message stating "Command not supported on this platform" is displayed.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--enable *feature*

Enables a feature on the switch. Valid values for *feature* include the following:

fcr

Enables the FC Routing service on the switch. The FC Routing service is disabled by default. Use **fosConfig --show** to determine if FCR is enabled or disabled.

When enabling FCR, you may encounter one of the following system messages:

"FC Routing service is in the process of being disabled, please try again after a few minutes." This means that a command to disable the FC Routing service is still in progress. Wait a few minutes and try again.

"FC Routing service is already enabled." This means that the FC Routing service is already enabled.

vf

Enables Virtual Fabrics. By default, Virtual Fabrics are enabled. Before enabling the Virtual Fabrics, ensure that there are no Administrative Domains (ADs) in effect. Administrative Domains must be disabled before enabling Virtual Fabrics. This command prompts for confirmation, because the chassis reboots after this command is executed.

ethsw

Enables the Ethernet switch service on the switch. The Ethernet switch service is disabled by default. Enabling the Ethernet switch service does not disrupt the FC traffic.

--disable *feature*

Disables a feature on the switch. Valid values for *feature* include the following:

fcr

Disables the FC Routing service on the switch. All enabled EX_Ports and VEX_Ports on the switch must be offline for this command to succeed.

To use this command to disable the FC Routing service only instead of disabling the switch, issue this command, then change the BB fabric ID using **fcrConfigure**.

When disabling the FCR service, you may encounter the following system messages:

"Please disable all EX/VEX_Ports first before running this command." This means that there were EX_Ports or VEX_Ports online when this command was issued. Take these ports offline and try the command again.

"FC Routing service is already disabled" -This means that the FC Routing service is already disabled.

vf

Disables Virtual Fabrics on the switch. This command prompts for confirmation, because the chassis reboots after this command is executed.

ethsw

Disables the Ethernet switch service. Ethernet switching is disabled by default. This operation reboots the switch and is therefore disruptive.

--show

Displays the current operating status of features on the switch.

Examples

To display the operating status of the services on a Brocade 6510:

```
switch:admin> fosconfig --show
FC Routing service:           disabled
Virtual Fabric:              enabled
Ethernet Switch Service:     Service not supported on this Platform
```

To disable the FC Routing service:

```
switch:admin> fosconfig --disable fcr
FC Routing service is disabled
```

To enable the FC Routing service:

```
switch:admin> fosconfig --enable fcr
FC Routing service is enabled
```

To enable Virtual Fabrics:

```
switch:admin> fosconfig --enable vf
WARNING: This is a disruptive operation that \
requires a reboot to take effect.
All EX ports will be disabled upon reboot.
Would you like to continue [Y/N]y
```

To disable Virtual Fabrics:

```
switch:admin> fosconfig --disable vf
WARNING: This is a disruptive operation that \
requires a reboot to take effect.
Would you like to continue [Y/N]y
```

To enable the Ethernet Switching service:

```
switch:admin> fosconfig --enable ethsw
WARNING: This operation will enable the \
Ethernet Switch Service on this switch.
Would you like to continue [Y/N]: y
Enabling the Ethernet Switch Service. Please wait ...
Create Ethernet switch instance 0
All service instances in sync.
The Ethernet Switch Service has been enabled.
Please enable all of the Blades with ID 74 now.
```

To disable the Ethernet Switching service:

```
switch:admin> fosconfig --disable ethsw
WARNING: This is a disruptive operation that \
requires a reboot to take effect.
Would you like to continue [Y/N]: y
The Ethernet Switch Service has been disabled.
Your system is being rebooted now.
The system is coming up, please wait...
```

SEE ALSO [fcrConfigure](#), [switchShow](#)

fosExec

Executes any Fabric OS command on a specified remote domain or logical switch context or on all remote domains or logical switches.

Synopsis

```
fosexec --fid FID -cmd "cmd [args]"
fosexec --fid all [-force] -cmd "cmd [args]"
fosexec --domain DID -cmd "cmd [args]"
fosexec --domain all -cmd "cmd [args]"
```

Description Use this command to execute any Fabric OS command on a specified remote domain or logical switch context or on all remote domains or logical switches.

The target logical switch is identified by its fabric ID (FID). When used with the **--fid all** option, the specified command is executed in all logical switches.

The target remote domain is identified by its Domain ID (DID). When used with the **--domain all** option, the specified command is executed in all remote domains.

The FIDs must be part of the FID permission list associated with the user account. If you execute **fosexec --fid all** and you do not have permission to one or more of the logical switches (FIDs), the command prompts for confirmation to continue with the execution of the command in the FIDs for which you do have permission. You can override the prompt for confirmation with the **-force** option. Refer to **userConfig** help for more information on logical switch access permissions.

Executing chassis-level commands through **fosexec --fid all** results in redundant output.

Use the **configure** command to set the Remote Fosexec feature. By default, the Remote Fosexec feature is disabled.

Notes The **--domain** option can be executed only if the Remote Fosexec feature is enabled on local and remote switches. The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|--|---|
| -cmd "<i>cmd</i> [<i>args</i>]" | Specifies the command to be executed including command options if applicable. The entire command string must be enclosed in double quotation marks. |
| --fid <i>FID</i> | Specifies the Fabric ID of the logical switch for which the command is executed. |
| --fid all | Executes the specified command in all logical switch contexts configured on the physical switch. |

- force** Executes the specified command on all logical switches in your permission list without prompting for confirmation.
- domain *DID*** Specifies the Domain ID of the remote switch for which the command is executed.
- domain all** Executes the specified command on all switches in the fabric.

Examples

To display switch information for a logical switch with FID 10:

```
switch:admin> fosexec --fid 10 -cmd "switchshow"
-----
"switchshow" on FID 10
switchName:      switch_10
switchType:      66.1
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:    1
switchId:        fffc01
switchWwn:       10:00:00:05:1e:41:5f:c2
zoning:          ON (lsan_cfg)
switchBeacon:    OFF
FC Router:       OFF
Allow XISL Use:  ON
LS Attributes:   [FID: 10, Base Switch: No, \
                  Default Switch: No, Address Mode 0]
Index Slot Port Address Media Speed State      Proto
=====
  0   1   0   640000  --   N8   No_Module  FC
  1   1   1   640100  --   N8   No_Module  FC
  2   1   2   640200  --   N8   No_Module  FC
```

To enable port 5/O on all logical switches:

```
switch:admin> fosexec --fid all -cmd "portenable 5/O"
-----

"portenable" on FID 128:
-----

"portenable" on FID 10:
A port or ports is/are not part of this switch.
-----

"portenable" on FID 20:
A port or ports is/are not part of this switch.
```

To display the firmware version for all logical switches:

```
switch:user> fosexec --fid all -cmd "firmwreshow -v"
LF permission does not exist for one or more \
  logical switches.
Would you like to continue [Y/N] : y
```

```
-----
"firmwreshow" on FID 10:
Appl      Primary/Secondary Versions
```

```
-----
FOS       v7.0.0
          v7.0.0
```

```
-----
"firmwreshow" on FID 20:
Appl      Primary/Secondary Versions
```

```
-----
FOS       v7.0.0
          v7.0.0
```

To display the switch name for all logical switches without confirmation:

```
switch:user> fosexec --fid all -force -cmd "switchname"
```

```
-----
"switchname" on FID 10:
```

```
switch_10
```

```
-----
"switchname" on FID 20:
```

```
switch_20
```

To display the fabric information for all switch in the fabric:

```
switch:user> fosexec --domain all -cmd "fabricshow"
```

```
Domain 3
```

```
=====
```

| Switch ID | Worldwide Name | Enet IP Addr | FC IP Addr | Name |
|-----------|--------------------------------|--------------|------------|---------|
| 3: | fffc03 10:00:00:05:1e:61:28:0a | 10.17.33.66 | 0.0.0.0 | "TOM-1" |
| 4: | fffc04 10:00:00:05:1e:61:28:22 | 10.17.33.68 | 0.0.0.0 | "TOM_4" |

```
The Fabric has 2 switches
```

```
Fabric Name: POD_1
```

```

Domain 4
=====
Switch ID   Worldwide Name           Enet IP Addr   FC IP Addr     Name
-----
   3: fffc03 10:00:00:05:1e:61:28:0a 10.17.33.66    0.0.0.0        >"TOM-1"
   4: fffc04 10:00:00:05:1e:61:28:22 10.17.33.68    0.0.0.0        "TOM_4"
The Fabric has 2 switches
Fabric Name: POD_1

```

To display the fabric information for domain 3 in the fabric:

```

switch:user> foexec --domain 3 -cmd "fabricshow"
Domain 3
=====
Switch ID   Worldwide Name           Enet IP Addr   FC
IP Addr     Name
-----
   3: fffc03 10:00:00:05:1e:61:28:0a 10.17.33.66    0.0.0.0        "TOM-1"
   4: fffc04 10:00:00:05:1e:61:28:22 10.17.33.68    0.0.0.0        "TOM_4"

The Fabric has 2 switches
Fabric Name: POD_1

```

SEE ALSO [setContext](#), [userConfig](#)

frameLog

Displays information about discarded frames.

Synopsis

```

framelog --disable [-type timeout | du | unroute]
framelog --enable [-type timeout | du | unroute]
framelog --clear
framelog --status
framelog --show [-type timeout | du | unroute]
  [-txport [slot]port]
  [-rxport [slot]port]
  [-sid source_PID][-did destination_PID]
  [-sfid fabric_ID] [-dfid fabric_ID>]
  [-mode summary | dump] [-n num_items]

```

Description Use this command to disable or re-enable the frame log, and to display detailed information about the discarded frames logged by the frame log. The frame log stores information about frames discarded due to certain reasons. The frame log sees information about only 40 discarded frames per second per chip on the chassis.

When frame drops occur on a switch, the frame log can help you identify to which flows the dropped frames belong and zero in on the affected applications by finding out the end-points of the dropped frame.

Use the **--show** option to display detailed information about the discarded frames that match the filter criteria provided by this command. The **--show** option supports two viewing modes: summary and dump mode. In summary mode, the command aggregates similar frames (those that have the same log timestamp, TX port, RX port, SID, DID, SFID and DFID) without displaying the raw frame contents. In dump mode, the command displays the raw frame contents, but the frames are not aggregated in the manner of the summary mode. In either mode, you can specify additional filters to customize the number and properties of the frames for closer examination.

When viewed in summary mode, the frame log displays the following information:

Log timestamp

The date and time when the frames described in this line of output were discarded. This value is accurate to within one second, that is, the displayed frames were not necessarily dropped at precisely the same time.

TX port

Egress port where the frames were bound to exit the chassis. A port displayed as -1 (or -1/-1 on slot-based systems) indicates an internal port.

RX port

Ingress port where the frames entered the chassis. A port displayed as -1 (or -1/-1 on slot-based systems) indicates an internal port.

SID

Source port ID in hexadecimal PID format.

DID

Destination port ID in hexadecimal PID format.

| | |
|--------------|---|
| SFID | Source Fabric ID |
| DFID | Destination Fabric ID |
| Type | Reason for the frame discard. Frames can be discarded for a variety of reasons. The discard reasons currently handled by the frame log are timeout, unroutable, and du. |
| Count | The number of discarded frames logged in the frame log that have the log timestamp, TX port, RX port, SID, DID, SFID and DFID values listed on this line of output. Note that this count may be less than the actual number of such frames discarded, because the frame log cannot log the details of all discarded frames. |

When viewed in dump mode, the frame log displays the following information:

| | |
|-----------------------|---|
| Log timestamp | The date and time when the frames described in this line of output were discarded. |
| TX port | Egress port where the frames were bound to exit the chassis. A port displayed as -1 (or -1/-1 on slot-based systems) indicates an internal port. |
| RX port | Ingress port where the frames were bound to enter the chassis. A port displayed as -1 (or -1/-1 on slot-based systems) indicates an internal port. |
| Type | Reason for the frame discard. Frames can be discarded for a variety of reasons. The discard reasons currently handled by the frame log are timeout, unroutable, and du. |
| Frame contents | The first 64 bytes of the frame contents in hexadecimal format. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The Frame Log operation is chassis-wide.

Operands This command has the following operands:

| | |
|--|--|
| --disable | Disables logging of discarded frames, which is enabled by default. This command clears the history of discarded frames and gathering of new information ceases. When frame logging is disabled, only the --help and --enable operations are permitted. |
| -type timeout du unroutable | Specifies the discard frame types to disable. Specify timeout for timeout discard frames, du for destination unreachable frames, and unroutable for unroutable frames. |
| --enable | Enables logging of discarded frames after the feature was disabled. |

- type timeout | du | unrout**
Specifies the discard frame types to enable. Specify timeout for timeout discard frames, du for destination unreachable frames, and unrout for unroutable frames.
- clear**
Clears the stored history of discarded frames. This operation is permitted only when the feature is enabled.
- status**
Displays the current status of the frame log service, and the type of discard frames that are being logged.
- show**
Prints a listing of details about stored frames that match the specified filter criteria. This operation is permitted only when the feature is enabled. You can specify one or more of the following options to filter the output. When used without operands, the **--show** command displays the unfiltered output. This option supports specifying that the TX port or RX port of displayed frames should be a backend port.
- type timeout | du | unrout**
Specifies the discard frame types to display. Specify timeout for timeout discard frames, du for destination unreachable frames, and unrout for unroutable frames.
- txport [slot]port**
Displays only the frames that were bound to exit by the specified egress port. The **-txport** option accepts argument "-1" (for fixed-port switches) or "-1/-1" (for modular switches). These stand for "any backend port". Using this notation, you can select specifically those discarded frames that have a backend port in the TX port field. Individual backend ports cannot be specified, only the quality of being a backend port can be specified.
- rxport [slot]port**
Displays only the frames that entered the chassis on the specified port. The **-rxport** option accepts argument "-1" (for fixed-port switches) or "-1/-1" (for modular switches). These stand for "any backend port". Using this notation, you can select specifically those discarded frames that have a backend port in the RX port field. Individual backend ports cannot be specified, only the quality of being a backend port can be specified.
- sid source_PID**
Displays only the frames with the specified 24-bit source address in the FC header. The port ID (PID) must be written in hexadecimal and must be prefixed by 0x.
- did destination_PID**
Displays only the frames with the specified 24-bit destination address in the FC header. The port ID (PID) must be written in hexadecimal and must be prefixed by 0x.
- sfid fabric_ID**
Displays only the frames with the specified fabric ID as the source fabric ID.
- sdfid fabric_ID**
Displays only the frames with the specified fabric ID as the destination fabric ID.

-mode summary | dump

Specifies the command output mode. In dump mode, the command prints the 64 bytes of frame contents. In summary mode, the frame contents are not displayed. Summary mode is the default.

-n num_items

Displays the specified number of items. In summary mode, an item is a record of frames with the same values of log timestamp, TX port, RX port, SID, DID, SFID and DFID. Note that more frames may be represented in the lines of output than specified in the *num_items* value, because one line can aggregate multiple frames in summary mode. In dump mode, each item represents a single frame. The default value of *num_items* is 20. The maximum value in summary mode is 1200, and in dump mode it is 12,000.

Examples

To display the status of the frame log service and the discard frame type:

```
switch:user> framelog --status
Service Status:           Enabled
Enabled Disc Frame Types:  timeout unroute
```

To display the framelog when unroute type is enabled:

```
switch:user> framelog --show
=====
=====
          Mon Mar 31 23:49:13 UTC 2014
=====
=====
Log      TX      RX
timestamp port  port  SID          DID          SFID  DFID  Type
Count
=====
=====
Mar 31 23:49:04 --    0    0x010000    0x0a1600    128   128   unroute
20
```

To display eight lines of an unfiltered frame log view in summary mode:

```
switch:user> framelog --show -mode summary -n 8
=====
=====
          Mon Jan 31 23:54:59 UTC 2011
=====
=====
Log      TX  RX
timestamp port port SID          DID          SFID  DFID  Type    Count
=====
=====
Jan 31 23:49:37 2   2   0xfffffd  0xfffffd   1     1   timeout  1
Jan 31 23:49:37 2   1   0x051500  0x060100   1     1   timeout  4
Jan 31 23:49:37 2   0   0x051700  0x060000   1     1   timeout  4
Jan 31 23:49:36 2   1   0x051500  0x060100   1     1   timeout  3
Jan 31 23:49:36 2   0   0x051700  0x060000   1     1   timeout  3
Jan 31 23:49:35 2   1   0x051500  0x060100   1     1   timeout  2
```

To display the frame contents in dump mode:

```
switch:user> framelog --show -mode dump -n 4
```

```
=====
                Mon Jan 31 18:34:47 UTC 2011
=====
Log           TX   RX
timestamp     port port Type   Frame contents (first 64 bytes)
=====
Jan 31 18:34:46 7/32 7/33 timeout
                00 03 a0 00 00 03 a1 00 00 28 00 00 00 00 00 00
                80 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Jan 31 18:34:46 7/32 7/33 timeout
                00 03 a0 00 00 03 a1 00 00 28 00 00 00 00 00 00
                80 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Jan 31 18:34:46 7/32 7/33 timeout
                00 03 a0 00 00 03 a1 00 00 28 00 00 00 00 00 00
                80 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Jan 31 18:34:46 7/32 7/33 timeout
                00 03 a0 00 00 03 a1 00 00 28 00 00 00 00 00 00
                80 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

SEE ALSO None

fspfShow

Displays Fabric Shortest Path First (FSPF) protocol information.

Synopsis `fspfshow`

Description Use this command to display FSPF protocol information and internal data structures of the FSPF module. The command displays the following fields:

| | |
|---------------------------|---|
| version | Version of the FSPF protocol. |
| domainID | Domain number of the local switch. |
| switchOnline | State of the local switch. |
| switchGenNum | Switch generation number. |
| switchGenNum(mmap) | Switch's current generation number managed by the switch driver. |
| domainValid | TRUE if the domain of the local switch is currently confirmed. |
| isl_ports | Bit map of all ISL ports. Bit positions correspond to the default areas of the ports. Bit 0 refers to default area of the switch, bit 1 refers to default area 1, and so forth. |
| trunk_ports | Bit map of all trunk slave ports. |
| f_ports | Bit map of all FX_Ports. |
| seg_ports | Bit map of all segmented ports. |
| active_ports | Bit map of all online ports. |
| minLSArrival | FSPF constant. |
| minLSInterval | FSPF constant. |
| LSoriginCount | Internal variable. |
| startTime | Start time of the FSPF task from boot time, in milliseconds. |

| | |
|----------------------|---|
| fspfQ | FSPF input message queue. |
| fabP | Pointer to fabric data structure. |
| agingTID | Aging timer ID. |
| agingTo | Aging time out value in milliseconds. |
| IsrDlyTID | Link State Record delay timer ID. |
| IsrDelayTo | Link State Record delay time out value in milliseconds. |
| IsrDelayCount | Counter of delayed Link State Records. |
| ddb_sem | FSPF semaphore ID. |
| event_sch | FSPF scheduled events bit map. |

Notes The output displays only the lines with a bit set. If a port bitmap does not have any bits set, the output displays as "None" for the first line of the bitmap.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display FSPF protocol information:

```
switch:admin> fspfshow
version                = 2
domainID               = 1
switchOnline           = TRUE
switchGenNum           = 4
switchGenNum(mmap)    = 4
domainValid            = TRUE
isl_ports              = None
trunk_ports            = None
fports                 = None
seg_ports              = None
active_ports           = None
le_ports               = None
nbStFullPortsP        = None
ve_portsP              = None
minLSArrival           = 3000
minLSInterval          = 5000
LSoriginCount          = 0
startTime              = 0
fspfQ                  = 0x111ffba8
```

```
fabP                = 0x11209bc8
agingTID            = 0x11329ae0
agingTo             = 10000
lsrDlyTID           = 0x1132ac90
lsrDelayTo          = 5000
lsrDelayCount       = 0
ddb_sem             = 0x11209bd8

fabP:
event_sch           = 0x0
max port            = 200
max phy port        = 40
```

SEE ALSO [bcastShow](#), [topologyShow](#), [uRouteShow](#)

gePortErrShow

Displays error statistics of Gigabit Ethernet (GbE) port and XGE ports.

Synopsis `geporterrshow`

`geporterrshow --help`

Description Use this command to display error statistics of Gigabit Ethernet (GbE) ports on the Brocade FX8-24 extension blade. Values for the following parameters are displayed:

frames tx Number of frames transmitted (Tx).

frames rx Number of frames received (Rx).

crc err Number of CRC errors.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

--help Displays command usage.

Examples To display error statistics:

```
switch:admin> geporterrshow
          frames      crc
          tx        rx   err
          =====
1/ge0   :    0      0    0
1/ge1   :    0      0    0
1/ge2   :    0      0    0
1/ge3   :    0      0    0
1/ge4   :    0      0    0
1/ge5   :    0      0    0
1/ge6   :    0      0    0
1/ge7   :    0      0    0
1/ge8   :    0      0    0
1/ge9   :    0      0    0
```

SEE ALSO [portCfgShow](#), [portErrShow](#)

h

Displays shell history.

Synopsis **h**

history

Description Use this command to view the shell history. The shell history mechanism is similar to the UNIX shell history facility. The **h** command displays the 20 most recent commands typed into the shell; the oldest commands are replaced as new ones are entered.

Operands None

Examples To display previous shell commands:

```
switch:admin> h
1 version
2 switchshow
3 portdisable 2
4 portenable 2
5 switchshow
```

SEE ALSO None

haDisable

haDisable

Disables the High Availability feature.

Synopsis **hadisable**

Description Use this command to disable the High Availability (HA) feature on a switch. If the HA feature is already disabled, this command does nothing.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the High Availability feature:

```
switch:admin> hadisable  
HA is disabled
```

SEE ALSO [haEnable](#), [haShow](#)

haDump

Displays High Availability status information.

Synopsis `hadump`

Description Use this command to display information about the status of the High Availability (HA) feature on a switch or a chassis. This command displays the following information:

- Time Stamp
- Local CP state (slot number and CP ID)
- Remote CP state (slot number and CP ID)
- Type of recovery (warm or cold)
- High Availability (enabled or disabled)
- Heartbeat (up or down)
- Health of standby CP
- HA synchronization status
- IP and Fibre Channel addresses configured for the switch or chassis.
- Additional internal HA state information, subject to change.

The health of the standby CP is defined as follows:

Healthy

The standby CP is running and the background health diagnostic has not detected any errors.

Failed

The standby CP is running, but the background health diagnostic has discovered a problem with the blade. Check the logs to determine an appropriate course of action. Failover is disabled until the standby CP is repaired. Information about the failing device in the standby CP is displayed.

Unknown

The standby CP health state is unknown because the standby CP does not exist, heartbeat is down, or Health Monitor detects a configuration file error.

The High Availability synchronization status is defined as follows:

HA State synchronized

The system is currently fully synchronized. If a failover becomes necessary, it is nondisruptive.

HA State not in sync

The system is unable to synchronize the two control processors (CPs) because the standby CP is faulty, an **haSyncStop** command was issued, or a system error occurred. If a failover becomes necessary at this time, active CP reboots and the failover is disruptive.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To view information about the High Availability feature status on a Brocade DCX 8510-8:

```
switch:admin> hadump
-----
TIME_STAMP: Mar 30 12:57:35.109173
-----
Local CP (Slot 7, CP1): Active, Warm Recovered
Remote CP (Slot 6, CP0): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized

CHASSIS
Ethernet IP Address: 10.20.114.57
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.20.114.58
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.20.112.1

CP1
Ethernet IP Address: 10.20.114.59
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.20.112.1

Backplane IP address of CP0 : 10.0.0.5
Backplane IP address of CP1 : 10.0.0.6
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
IPv6 Gateways:
-----
TIME_STAMP: Mar 30 12:57:37.31282
-----
== Service ==
chassis0:0(8.0)
major:          8      part:          0
lo:             0      role:         2
id:             0      dump:        3
epoch:         3      log:         0
block:         0      bits:       321
recov:         0      index:      14
upper:        14      coldnum:    3
warmnum:     14      stbynum:   4
stat:         4      prevstat:  4
sv_winsz:    64      nscm:     10
== Config ==
(Output truncated)
```

SEE ALSO [haFailover](#), [haShow](#)

haEnable

Enables the High Availability feature.

Synopsis `haenable`

Description Use this command to enable the High Availability (HA) feature on a switch. If the HA feature is already enabled, this command has no effect.

If the HA feature is disabled, this command enables it. The standby CP reboots as part of the process. The command displays a warning message and prompts for confirmation before rebooting the CP.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable the High Availability feature on a Brocade DCX 8510-8:

```
switch:admin> haenable
Warning: This command will enable the HA. It will reboot the standby
CP and require all telnet, secure telnet, and SSH sessions to the
standby CP to be restarted

Are you sure you want to go ahead [y/n]?y
y
```

To verify that High Availability is enabled:

```
switch:admin> hashow
Local CP (Slot 7, CP1): Active, Warm Recovered
Remote CP (Slot 6, CP0): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
```

SEE ALSO [haDisable](#), [reboot](#)

haFailover

Forces the failover mechanism so that the standby control processor (CP) becomes the active CP.

Synopsis `hafailover`

Description Use this command to force the failover mechanism to occur so that the standby CP becomes the active CP. In case the active and standby CPs are not synchronized or the system is not in redundant mode, the command aborts.

Notes When High Availability (HA) synchronization is enabled and the CPs are in sync, the port traffic light does not flash during the failover, even while traffic is continuing to flow.

This command is supported only on dual-CP systems.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To force the failover of the active CP to the standby CP in the switch:

```
switch:admin> hafailover
Local CP (Slot 7, CP1): Active, Warm Recovered
Remote CP (Slot 6, CP0): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
Warning: This command is being run on a redundant control processor(CP)
system, and this operation will cause the active CP to reset.
Therefore all existing telnet sessions are required to be restarted.

Are you sure you want to fail over to the standby CP [y/n]?
```

SEE ALSO [haDisable](#), [haEnable](#), [haShow](#)

haReboot

Performs high availability (HA) reboot.

Synopsis `hareboot`

Description Use this command to perform warm reboot in the switch-based systems. On chassis-based systems, this command behaves similar to the **haFailover** command so that the standby CP becomes the active CP.

In chassis-based systems, if the active and standby CPs are not synchronized or the system is not in the redundant mode, the command aborts.

When the switch-based system is not fully up or stable, the **haReboot** command cannot be completed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To perform Ha reboot in the switch-based systems:

```
switch:admin> hareboot
HA Rebooting ...
```

To perform Ha reboot in the chassis-based systems:

```
switch:admin> hareboot
Restarting system with command 'Machine Rebooting
```

SEE ALSO [haDisable](#), [haEnable](#), [haShow](#)

haRedundancy

Displays switch uptime.

Synopsis **haredundancy --show**

haredundancy --help

Description Use this command to display the switch uptime and the chassis control processor redundancy statistics. For chassis, the command displays the control processor redundancy settings and switch uptime. For switches, the command displays the switch uptime only.

The display includes the following information:

- Current active session: Displays the settings for the current session.
 - HA synchronization status.
 - Active slot state: Displays CP ID, whether CP is local or remote, and recovery type information.
 - Standby slot state: Displays CP ID, and whether CP is local or remote.
 - Start time: Displays the start time of the services in sync state.
- Previous active session: Displays the settings for the previous active session.
 - Active slot state: Displays CP ID, and recovery type information.
 - Standby slot state: Displays CP ID information.
 - Start time: Displays the start time of the services in sync state for the session.
 - End time: Displays end time of the session caused by expected or unexpected recovery.
- System uptime: Displays the start time of the system services. This changes when there is a power cycle or both CPs are reset together.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--show Displays the switch uptime and other information.

--help Displays the command usage.

Examples To display switch uptime with a health standby CP:

```
switch:admin> haredundancy --show
=== HA Redundancy Statistics ===
      HA State synchronized
      Current Active Session:
      Active Slot = CP0 (Local), Expected Recovered
      Standby Slot = CP1 (Remote)
      Start Time: 17:55:33 UTC Fri Jan 03 2014

      Previous Active Session:
      Active Slot = CP1, Expected Recovered
```

```
Standby Slot = CP0
Start Time: 17:49:46 UTC Fri Jan 03 2014
End Time: 17:54:10 UTC Fri Jan 03 2014

System Uptime: 17:42:11 UTC Fri Jan 03 2014
```

To display switch uptime without a standby CP:

```
switch:admin> haredundancy --show
=== HA Redundancy Statistics ===
Non-redundant

System Uptime: 17:42:11 UTC Fri Jan 03 2014
```

To display switch uptime:

```
switch:admin> haredundancy --show
=== HA Redundancy Statistics ===

Not supported on this platform

System Uptime: 09:42:12 UTC Wed Nov 22 2013
```

SEE ALSO [haShow](#)

haShow

Displays control processor (CP) status.

Synopsis `hashow`

Description Use this command to display the control processor status. The display includes the following information:

- Local CP state (slot number and CP ID), warm or cold, recovering or recovered.
- Remote CP state (slot number and CP ID).
- High Availability (enabled or disabled).
- Heartbeat (up or down).
- Health of standby CP
- HA synchronization status

The health of the standby CP is defined as follows:

Healthy

The standby CP is running and the background health diagnostic has not detected any errors.

Failed

The standby CP is running, but the background health diagnostic has discovered a problem with the blade. Check the logs to determine the appropriate action. Failover is disabled until the standby CP is repaired. Information about the failing device in the standby CP is displayed.

Unknown

The standby CP health state is unknown because of one of the following reasons: the standby CP does not exist, Heartbeat is down, or the Health Monitor has detected a configuration file error.

The High Availability synchronization status is defined as follows:

HA State synchronized

The system is currently fully synchronized. If a failover becomes necessary, it is nondisruptive.

HA State not in sync

The system is unable to synchronize the two CPs. This may be caused by one or more of the following conditions:

- The **haFailover** command was issued. In this case the "HA State not in sync" state is transitory.
- The standby CP is faulty.
- The **haSyncStop** command was issued.
- A system error occurred.

If a failover becomes necessary while the CPs are not in sync, the standby CP reboots, and the failover is disruptive.

Notes This command may not be supported on nonbladed systems.

Slot numbers for CP1 and CP0 vary depending on the hardware platform. On the Brocade DCX 8510-8, CP0 is in slot 6 and CP1 is in slot 7. On the Brocade DCX 8510-4, CP0 is in slot 4 and CP1 is in slot 5. On the Brocade X6 Director, CP0 is in slot 1 and CP1 is in slot 2.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display CP status on a Brocade DCX 8510-8 with a healthy standby CP:

```
switch:admin> hashow
Local CP (Slot 7, CP1) : Active, Warm Recovered
Remote CP (Slot 6, CP0) : Standby, Healthy
HA Enabled, Heartbeat Up, HA State Synchronized
```

To display CP status on a Brocade X6 Director with a healthy standby CP:

```
switch:admin> hashow
Local CP (Slot 1, CP0): Active, Cold Recovered
Remote CP (Slot 2, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
```

SEE ALSO [haDisable](#), [haEnable](#), [haFailover](#)

haSyncStart

Enables High Availability state synchronization.

Synopsis `hasyncstart`

Description Use this command to enable the High Availability (HA) state synchronization.

After issuing **haSyncStop**, the switch does not go back to sync start unless you perform one of the following actions: reboot the active CP, reboot the standby CP, insert a new standby CP (blade), or issue the **haSyncStart** command. The time it takes for the HA sync to complete depends on the system configuration. The HA sync timeout is set to 10 minutes.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Examples To enable the HA state synchronization:

```
switch:admin> hasyncstart
```

SEE ALSO [haFailover](#), [haShow](#), [haSyncStop](#)

haSyncStop

Disables High Availability state synchronization.

Synopsis `hasyncstop`

Description Use this command to temporarily disable High Availability (HA) synchronization.

After issuing **haSyncStop**, the switch does not go back to sync start unless you perform one of the following actions: reboot the active CP, reboot the standby CP, insert a new standby CP (blade), or issue the **haSyncStart** command. The time it takes for the HA sync to complete depends on the system configuration. The HA sync timeout is set to 10 minutes.

Notes Disabling HA synchronization may cause failover to be disruptive.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the HA state synchronizing process:

```
switch:admin> hasyncstop
```

SEE ALSO [haFailover](#), [haShow](#), [haSyncStart](#)

help

Displays command help information.

Synopsis `help [command]`

`help [-p | -page]`

`help --help`

Description Use this command without operands to display an alphabetical listing of commands for which help is available. When used without an operand, the command listing displays without page break.

Pipe the output through the Unix **grep** command to filter the output.

Use the **-page** operand to display the commands for which help is available one page at a time. Press **Enter** to go to the next page. When using help with the **-page** option, you can search for specific strings by entering a forward slash, followed by a text string, for example **/zone**.

The help listing includes only commands that are available to the current user; command availability may vary depending on the following conditions:

- Login user role
- License key
- Hardware platform

To access help information for a specific command, enter the command name as an operand.

Commands ending in "Help" display grouped commands for a particular subsystem; for example, the **diagHelp** command displays a list of diagnostic commands.

Operands This command has the following operands:

command Specifies the name of the command for which to display help information. This operand is optional.

p | -page Displays help output with page breaks.

--help Displays the command usage.

Examples To display a listing of commands for which help is available (with page breaks):

```
switch:admin> help -p
aaaconfig      Configure RADIUS for AAA services
ad             Specifies all administration domain
              (AD)-level operations
ag            Configure the Access Gateway feature
agshow        Displays the Access Gateway information
              registered with the fabric
aliadd        Add a member to a zone alias
alcreate      Create a zone alias
```

```

alidelete      Delete a zone alias
aliremove      Remove a member from a zone alias
alishow        Print zone alias information
aptpolicy      Get and set Advanced Performance
                Tuning policy
auditcfg       Modifies and displays audit log filter
                configuration.
auditdump      Display audit log
                (output truncated)

```

To search for the string "zone" while paging is enabled (enter /zone after issuing the command):

```

switch:admin> help -p
aaaconfig      Configure RADIUS for AAA services
ad             Specifies all administration domain
                (AD)-level operations
ag            Configure the Access Gateway feature
agshow        Displays the Access Gateway information
                registere with the fabric
aliadd        Add a member to a zone alias
alcreate      Create a zone alias
alidelete     Delete a zone alias
aliremove     Remove a member from a zone alias
alishow       Print zone alias information
aptpolicy     Get and set Advanced Performance Tuning
                policy
auditcfg      Modifies and displays audit log filter
                configuration.
auditdump     Display audit log
authutil      Get and set configuration
backplanetest Backplane connection test for multi-blade
                systems.
backport      Test for back-end ASIC pair to ASIC pair links.
bannerset     Set security banner
bannershow    Display security banner
bcastshow     Display broadcast routing information
Type <CR> or <SPACE BAR> to continue, <q> to stop
/zone
aliadd        Add a member to a zone alias
alcreate      Create a zone alias
alidelete     Delete a zone alias
aliremove     Remove a member from a zone alias
alishow       Print zone alias information
aptpolicy     Get and set Advanced Performance Tuning policy
auditcfg      Modifies and displays audit log filter
                configuration.
                (output truncated)

```

To filter the output with the **grep** command:

```

switch:admin> help | grep errshow
errshow        Print error log
porterrshow    Print port summary
switch:admin>

```

SEE ALSO [diagHelp](#), [routeHelp](#), [zoneHelp](#)

historyLastShow

Displays the latest entry in the field replaceable unit (FRU) history log.

Synopsis `historylastshow`

Description Use this command to display the latest entry of the history log, which records insertion and removal events for field-replaceable units (FRUs), such as blades, power supplies, fans, and world wide name (WWN) cards. The type of FRU supported depends on the hardware platform.

Each history record contains three lines of information. The first line of each record contains the following fields:

Object type

On standalone platforms: FAN, POWER SUPPLY, WWN (WWN card), or UNKNOWN. On enterprise-class platforms: FAN, POWER SUPPLY, CORE BLADE (core switch blade), AP BLADE (application processor), SW BLADE (port blade), CP BLADE (control processor), WWN (WWN card), or UNKNOWN.

Object number

Slot number for blades. Unit number for all other object types.

Event type

Inserted, Removed, or Invalid.

Time of the event

Format: *Day Month dd hh:mm:ss yyyy*,

The second and third lines of a record contain the factory part number and factory serial number, if applicable.

Factory Part Number

xx-yyyyyyy-zz or Not available.

Factory Serial Number

xxxxxxxxxxxx or Not available.

The size of the history log depends on the hardware platform. The Brocade DCX supports a maximum log size of 50 entries. Because the WWNs are completely separate FRUs in the DCX, they have redundant data in these units. All other platforms that contain FRUs support 28 history log entries.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the late FRU insertion or removal event:

```
switch:admin> historylastshow
```

```
POWER SUPPLY Unit 2      Inserted at Tue Aug 14 15:52:10 2001  
Factory Part Number:    60-0001536-02  
Factory Serial Number:  1013456800
```

```
Records: 11
```

SEE ALSO [historyShow](#)

historyMode

Displays the mode of the field replaceable unit (FRU) history log.

Synopsis `historymode`

Description Use this command to display the mode of the history buffer, which records the insertion and removal of FRUs on a switch or chassis.

This command supports two modes of handling new log entries once the history buffer has reached its maximum size:

Rotating mode

Any new entry exceeding the maximum buffer size overwrites the oldest entry in the log. This is the default mode.

First-in mode

Any new entry exceeding the maximum buffer size is discarded. The original entries in the buffer is preserved.

The history mode is a factory setting that cannot be modified. The size of the history buffer depends on the hardware platform. The Brocade DCX backbone supports a maximum log size of 50 entries. Because the WWNs are completely separate FRUs in the DCX, they have redundant data in these units. All other platforms containing FRUs support 28 history log entries.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the mode of the history log:

```
switch:admin> historymode  
  
History Mode is: Rotating.
```

SEE ALSO [historyLastShow](#), [historyShow](#)

historyShow

Displays the entire field replaceable unit (FRU) history log.

Synopsis `historyshow`

Description Use this command to display the entire history log, which records insertion and removal events for field-replaceable units (FRUs), such as blades, power supplies, fans, and world wide name (WWN) cards. The type of FRU supported depends on the hardware platform.

Each history record contains three lines of information. The first line of each record contains the following:

Object type

On standalone platforms: FAN, POWER SUPPLY, WWN (WWN card), or UNKNOWN.

On enterprise-class platforms: FAN, POWER SUPPLY, CORE BLADE (core switch blade), SW BLADE (port blade), AP BLADE (application processor), CP BLADE (control processor), WWN (WWN card), or UNKNOWN.

Object number

Slot number for blades. Unit number for all other object types.

Event type

Inserted, Removed, or Invalid

Time of the event

Format: *Day Month dd hh:mm:ss yyyy*

The second and third lines of a record contain the factory part number and factory serial number, if applicable:

Factory Part Number

xx-yyyyyyy-zz or Not available

Factory Serial Number

xxxxxxxxxxxx or Not available

The size of the history buffer depends on the HW platform. The Brocade DCX supports a maximum log size of 50 entries. Because the WWNs are completely separate FRUs in the DCX, they have redundant data in these units. All other platforms that contain FRUs support 28 history log entries.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the history log of FRUs and removal events on a standalone switch:

```
switch:admin> historyshow

FAN Unit 3                Removed at Tue Aug 14 10:05:37 1970
Factory Part Number:      20-123456-12
Factory Serial Number:    1013456800

POWER SUPPLY Unit 1       Inserted at Tue Aug 14 10:52:10 1970
```

historyShow

```
Factory Part Number: 60-0001536-02
Factory Serial Number: Not Available

FAN Unit 3          Inserted at Tue Aug 14 10:23:45 2001
Factory Part Number: 20-123456-12
Factory Serial Number: 1013456800

WWN Unit 1          Inserted at Tue Aug 14 11:03:45 2001
Factory Part Number: 40-0000031-03
Factory Serial Number: 1013456800

SW BLADE Slot 3     Removed at Tue Aug 14 12:10:09 2001
Factory Part Number: 60-0001532-03
Factory Serial Number: 1013456800

CP BLADE Slot 6     Removed at Tue Aug 14 13:45:07 2001
Factory Part Number: 60-0001604-02
Factory Serial Number: FP00X600128

SW BLADE Slot 3     Inserted at Tue Aug 14 13:53:40 2001
Factory Part Number: 60-0001532-03
Factory Serial Number: 1013456800

CP BLADE Slot 6     Inserted at Tue Aug 14 13:59:50 2001
Factory Part Number: 60-0001604-02
Factory Serial Number: FP00X600128

POWER SUPPLY Unit 2 Inserted at Tue Aug 14 15:52:10 2001
Factory Part Number: 60-0001536-02
Factory Serial Number: 1013456800

Records: 11
```

SEE ALSO [historyLastShow](#)

Displays a process summary.

Synopsis `i [process/D]`

Description Use this command to display information about a specified process or about all processes running on the local switch. One line is displayed per process. Fields displayed with this command include the

| | |
|-------------------|---------------------------------------|
| F | Process flags: |
| ALIGNWARN | 001 print alignment warning messages |
| STARTING | 002 being created |
| EXITING | 004 getting shut down |
| PTRACED | 010 set if ptrace (O) has been called |
| TRACESYS | 020 tracing system calls |
| FORKNOEXEC | 040 forked but did not exec |
| SUPERPRIV | 100 used super-user privileges |
| DUMPCORE | 200 dumped core |
| SIGNALLED | 400 killed by a signal |
| S | Process state codes: |
| D | Uninterruptable sleep (usually IO) |
| R | Runnable (on run queue) |
| S | Sleeping |
| T | Traced or stopped |
| Z | A defunct ("zombie") process |

| | |
|--------------|--|
| UID | The effective user ID number of the process |
| PID | The process ID of the process |
| PPID | The process ID of the parent process |
| C | Processor utilization for scheduling |
| PRI | Priority number of the process; higher numbers mean lower priority |
| NI | Nice value used in priority computation |
| ADDR | Memory address of the process |
| SZ | The total size of the process in virtual memory, in pages |
| WCHAN | The address of an event for which a process is sleeping (if blank, process is running) |
| TTY | The controlling terminal of the process (? displayed for no controlling terminal) |
| TIME | The cumulative execution time for the process |
| CMD | The command name of the process. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

processID Specifies the process name or process ID for the process to display.

Examples To display information about process ID 433:

```
switch:admin> i 433
F  S UID PID PPID C  PRI NI ADDR  SZ  WCHAN TTY  TIME  CMD
000 S 0 433  1 0  69  0  - 1283  5c64 ?  00:00:02 fabricd
```

SEE ALSO [diagHelp](#), [routeHelp](#)

iflShow

Displays the interfabric link (IFL) information.

Synopsis `iflshow`

`iflshow --help`

Description Use this command to display the current connection and information of the interfabric links (IFL) on an edge switch. The command output includes the following information:

| | |
|-----------------|---|
| E-Port | Port number of the local switch to which FC Router switch is connected |
| Ex-Port | Port Number of FC Router switch |
| FCR WWN | WWN of the FC Router switch |
| FCR FID | Fabric ID of FC Router switch |
| FCR Name | Switch name of FC Router |
| Speed | IFL connection speed, if applicable. Connection speed not applicable to VE_Ports. For these port types, speed displays as '-'. <ul style="list-style-type: none"> • TRUNK - Trunking enabled, if applicable • QOS - QoS enabled, if applicable • ENCRYPT - Encryption enabled, if applicable • COMPRESS - Compression enabled, if applicable • CR_RECOV - Credit recovery enabled, if applicable • FEC - Forward Error Correction enabled, if applicable |
| BW | Bandwidth of the link. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

`--help` Displays the command usage.

iflShow

Examples To display the interfabric links:

```
switch:admin> iflshow
  E-   EX-       FCR WWN           FCR  FCR  Speed  BW
  #   Port Port           FID  Name
-----
1: 7-> 8   10:00:00:05:1e:8b:68:10  8   fcr1  4G     4G TRUNK
2: 12-> 13  10:00:00:05:1e:a1:e3:3a  15  fcr2  4G     4G TRUNK
3: 15-> 15  10:00:00:05:1e:a1:e1:f9  12  sw87  4G     4G TRUNK QOS
4: 16-> 16  10:00:50:eb:1a:00:00:02  4   fcr3  -      4G
```

SEE ALSO [fcrEdgeShow](#), [fcrFabricShow](#)

interfaceShow

Displays FSPF interface information.

Synopsis `interfaceshow [slot][port]`

Description Use this command to display the two data structures associated with FSPF interfaces (E_Ports) on the switch:

- The permanently allocated interface descriptor block (IDB).
- The neighbor data structure. This data structure is allocated when a switch port becomes an E_Port. The neighbor data structure contains all the information relating to the switch that is connected to an adjacent switch.

This command displays the content of both data structures, if they have been allocated.

The following fields are displayed:

| | |
|--------------------|---|
| idbP | Pointer to IDB. |
| nghbP | Pointer to neighbor data structure. |
| ifNo | Interface number. |
| masterPort | Port number of the trunk master port, if present, of the trunk group of which this port is a part. |
| prgm_if_id | Internal variable |
| curr_if_id | Internal variable |
| prgmCnt | Number of remote domains the link has been programmed to reach for routing traffic. |
| defaultCost | Default cost of sending a frame over the interswitch link (ISL) connection to this interface. |
| cost | Cost of sending a frame over the ISL connected to this interface. A value of 1000 indicates a 1-Gbps link. A value of 500 indicates a 2-Gbps link. For links with a bandwidth greater than 2 Gbps, the cost is 500. For links with less than 1 Gbps, the cost is 2000. Refer to linkCost for more information. |
| ifBw | The rounded bandwidth of the interface, in Gbps. |
| totalBw | The summed total bandwidth for the entire trunk and ISL, in Gbps. |

interfaceShow

| | |
|-----------------------|--|
| delay | Conventional delay incurred by a frame transmitted on this ISL. A fixed value required by the FSPF protocol. |
| lastScn | Type of the last State Change Notification received on this interface. |
| lastScnTime | Time the last State Change Notification was received on this interface. |
| upCount | Number of times this interface came up, with respect to FSPF. |
| lastUpTime | Last time this interface came up. |
| downCount | Number of times this interface went down. |
| lastDownTime | Last time this interface went down. |
| downReason | Type of last State Change Notification that caused this interface to go down. |
| iState | Current state of this interface. The state can be UP or DOWN. An interface in DOWN state does not have an allocated neighbor data structure and cannot be used to route traffic to other switches. |
| state | Current state of this interface. This E_Port is used to route traffic to other switches only if the state is NB_ST_FULL. |
| lastTransition | Time the last state changed on this interface. |
| nghbCap | Neighbor capabilities. Should be 0. |
| nghbld | Domain ID of the neighbor (adjacent) switch. |
| idbNo | IDB number. Should be equal to <i>port</i> . |
| remPort | Port number on the remote switch connected to this port. |
| nflags | Internal FSPF flags. |
| initCount | Number of times this neighbor was initialized without the interface going down. |
| lastInit | Time of the last initializing state, NB_ST_INIT, on this interface. |

| | |
|----------------------------|---|
| firstHlo | Time of the first hello sent on this interface. |
| nbstFull | Time of the last finishing state, NB_ST_FULL, on this interface. |
| &dbRetransList | Pointer to the database retransmission list. |
| &lSrRetransList | Pointer to the link state records (LSR) retransmission list. |
| &lSrAckList | Pointer to the link state acknowledgements (LSA) retransmission list. |
| inactTID | Inactivity timer ID. |
| helloTID | Hello timer ID. |
| dbRtxTID | Database retransmission timer ID. |
| lSrRtxTID | LSR retransmission timer ID. |
| inactTo | Inactivity timeout value, in milliseconds. When this timeout expires, the adjacency with the neighbor switch is broken and new paths are computed to all possible destination switches in the fabric. |
| helloTo | Hello timeout value, in milliseconds. When this timeout expires, a Hello frame is sent to the neighbor switch through this port. |
| rXmitTo | Retransmission timeout value, in milliseconds. It is used to transmit topology information to the neighbor switch. If no acknowledgement is received within this value, the frame is retransmitted. |
| nCmdAcc | Total number of commands accepted from the neighbor switch. Number includes Hellos, Link State Updates (LSUs), and LSAs. |
| nInvCmd | Number of invalid commands received from the neighbor switch. Usually commands with an FSPF version number higher than the one running on the local switch. |
| nHloIn | Number of Hello frames received from the neighbor switch. |
| nInvHlo | Number of invalid Hello frames (Hello frames with invalid parameters) received from the neighbor switch. |

interfaceShow

| | |
|---------------------------|--|
| nLsuIn | Number of LSUs received from the neighbor switch. |
| nLsaIn | Number of LSAs received from the neighbor switch. |
| attHloOut | Number of attempted transmissions of Hello frames to the neighbor switch. |
| nHloOut | Number of Hello frames transmitted to the neighbor switch. |
| attLsuOut | Number of attempted transmissions of LSUs to the neighbor switch. |
| nLsuOut | Number of LSUs transmitted to the neighbor switch. |
| attLsaOut | Number of attempted transmissions of LSAs to the neighbor switch. |
| nLsaOut | Number of LSAs transmitted to the neighbor switch. |
| StuckCnt | Number of HLO timeouts that occurred before the port changed to the NB_ST_FULL state. |
| state | Substate of the port. The port can be in one of the following 12 substates: |
| INIT (0) | The port is initializing. |
| ROUTABLE (1) | The port is ready to receive frames. |
| ROUTABLE_WAIT (2) | The port is waiting for notification from the neighbor. |
| ROUTABLE_SEND (3) | The port is preparing for local route update; negotiates for the necessary locks before updating the routes. |
| ROUTABLE_TIMER (4) | The request for one or more locks failed; delay for a short interval before retrying. |
| DONE (5) | The port is online and in use. |
| DECOM_START (6) | The decommissioning request is sent to the neighbor and waiting for a response. |
| DECOM_WAIT (7) | Waiting for neighbor to remove routes using the interswitch link (ISL). |

| | |
|---------------------------|--|
| DECOM_SEND (8) | Preparing to update routes of the local port; negotiates for the necessary before updating the routes. |
| DECOM_TIMER (9) | The request for one or more locks failed; delay for a short interval before retrying. |
| DECOM_DONE (10) | The local routes are updated; waiting for acknowledgement from the neighbor. |
| DECOM_BLOCK (11) | The local routes are updated; the neighbor sent acknowledgement. |
| chassis_lock | Internal variable |
| lock_request | Internal variable |
| nbr_lock_lock | Internal variable |
| r_rdy_rcvd | Indicates if the neighbor has reported as ready to receive frames. |
| nbr r_rdy flags | Internal flags reported by the neighbor. |
| lock_busy_cnt | Number of times a lock has reported as busy. |
| decom_active | Indicates if a decommissioning request is active on the port. |
| decom_initiator | Indicates if the local port is the decommissioning request initiator. |
| decom_active_port | Indicates the port for which the decommissioning request is active. |
| decom_trunk_member | Indicates if decommissioning request is for a multi-link trunk. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|-------------|---|
| <i>slot</i> | For bladed systems only, specify the slot number of the port to be displayed, followed by a slash (/). |
| <i>port</i> | Specify the number of the port to be displayed, relative to its slot for bladed systems. Use switchShow for a list of valid ports. This operand is optional; if omitted, the interface information for all ports is displayed. |

When invoked without operands, this command displays the interface information for all ports on the switch (including non-E_Ports).

Examples To display FSPF interface information:

```
switch:user> interfaceshow 1/19
idbP          = 0x102651a0

Interface 131 data structure:

nghbP        = 0x10b968f0
ifNo         = 131
masterPort   = 131
prgm_if_id   = 0x43120021
curr_if_id   = 0x43120021
prgmCnt      = 1
defaultCost  = 500
cost         = 500
ifBw        = 8G
totalBw     = 24G
delay       = 1
lastScn     = 0
lastScnTime = Jan 01 00:00:00.000
upCount     = 0
lastUpTime  = Jan 01 00:00:00.000
downCount   = 0
lastDownTime = Jan 01 00:00:00.000
downReason  = 0
iState      = UP

Neighbor 131 data structure:

state        = NB_ST_FULL
lastTransition = Jun 04 05:05:02.592
nghbCap      = 0x0
nghbId       = 188
idbNo        = 131
remPort      = 243
nflags       = 0xf
initCount    = 0
lastInit     = Jun 04 05:05:02.592
firstHlo     = Jun 04 05:05:21.474
nbstFull     = Jan 01 00:00:00.000
&dbRetransList = 0x10b96928
&lsrRetransList = 0x10b96930
&lsrAckList   = 0x10b96938
inactTID     = 0x10b96b08
helloTID     = 0x10b97cb8
dbRtxTID     = 0x10b98e68
lsrRtxTID    = 0x10b9b1c8
inactTo      = 160000
helloTo      = 20000
rXmitTo      = 5000
nCmdAcc      = 38
nInvCmd      = 0
nHloIn       = 35
nInvHlo      = 0
nLsuIn       = 2
nLsaIn       = 1
```

```
attHloOut      = 35
nHloOut        = 35
attLsuOut      = 1
nLsuOut        = 1
attLsaOut      = 2
nLsaOut        = 2
StuckCnt       = 0
```

```
Port Sub-State Data:
state          = DONE (5)
chassis_lock   = 0
lock_request   = 0
nbr_lock_lock  = 0
r_rdy_rcvd    = 0
nbr_r_rdy flags = 0x0
lock_busy_cnt  = 0
decom_active   = 0
decom_initiator = 0
decom_active_port = -1
decom_trunk_member = 0
```

SEE ALSO [nbrStateShow](#), [portShow](#), [switchShow](#)

iodReset

Disables in-order delivery (IOD) on a switch.

Synopsis `iodreset`

Description Use this command to disable in-order delivery enforcement on the local switch. IOD is disabled by default, and can only be disabled after it has been enabled with the **iodSet** command. This command disables the legacy IOD enforcement only.

Disabling IOD allows faster rerouting after a fabric topology change, but it may cause out-of-order delivery of frames during fabric topology changes.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable IOD enforcement:

```
switch:admin> iodreset
```

```
IOD is not set
```

SEE ALSO [iodSet](#), [iodShow](#)

iodSet

Enables in-order delivery (IOD).

Synopsis `iodset`

`iodset --help`

Description Use this command to enforce in-order delivery of frames during a fabric topology change.

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for instance, a link goes down), traffic is rerouted around the failure and some frames might be delivered out of order. This command ensures that frames are not delivered out-of-order, even during fabric topology changes. It enforces a sufficient delay between the event that causes an existing path to be removed and the establishment of a new path, so that frames are delivered in order. However, this also means that frames are dropped during the delay, causing I/O failures.

When used without operands, **iodSet** enables in-order-delivery of frames on a switch (legacy IOD behavior). Frame loss is unavoidable when a port goes down.

IOD is disabled by default. Use **iodShow** to display current settings. Use **iodReSet** to restore the default setting.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--help
Displays the command usage.

Examples To display the default legacy IOD setting:

```
switch:admin> iodshow
IOD is not set
```

To enable IOD:

```
switch:admin> iodset
IOD is set
```

SEE ALSO [iodShow](#), [iodReset](#)

iodShow

Displays the in-order delivery (IOD) setting.

Synopsis `iodshow`

Description Use this command to display the IOD setting on the switch. By default, IOD is disabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the current setting of the IOD setting:

```
switch:admin> iodshow
```

```
IOD is not set
```

SEE ALSO [iodReset](#), [iodSet](#)

ipAddrSet

Sets the Ethernet and FC IP addresses.

Synopsis

```
ipaddrset [-ipv6 ] [--add x:x:x:x:x/n | --delete]

ipaddrset [-cp cp_number | -chassis]
  [-ipv6 ] [--add x:x:x:x:x/n | --delete]

ipaddrset [-slot slot] [-eth0 | -eth1] [--add x.x.x.x/n | --delete]

ipaddrset [-slot slot] -gate [--add x.x.x.x | --delete]

ipaddrset -ls FID --add IPv4_address/prefix

ipaddrset -ls FID --delete

ipaddrset -ipv6 [-auto | -noauto]

ipaddrset [-cp cp_number | -chassis]

ipaddrset -cp cp_number -ipv4 -add -host hostname
  -ethip ipaddress -ethmask mask
  -gwyip gateway_ip

ipaddrset -ipv4 -add -ethip ipaddress -ethmask mask
  -gwyip gateway_ip -dhcp [ON | OFF]

ipaddrset -chassis -ipv4 -add
  -ethip ipaddress -ethmask mask

ipaddrset -ipv6 [-dhcpv6 | -nodhcpv6]

ipaddrset --clear

ipaddrset --help
```

Description

Use this command to configure the IP addresses on a switch, a control processor (CP), a blade processor (BP), or a standalone application processor (AP). On platforms that support Logical Fabrics, this command configures the IPv4 Fibre Channel addresses for the logical fabric IPFC network interface. The IPFC (IP over Fibre Channel) protocol allows switches to send IP traffic over Fibre Channel rather than through Ethernet connections.

This command supports an interactive legacy mode and a command line interface. Use the command line interface to configure IPv6 addresses, to enable or disable stateless IPv6 autoconfiguration, to assign a Fibre Channel IPv4 address and prefix to a logical switch IPFC network interface, and to configure a service port on the CP blade of a Brocade DCX. When run interactively in legacy mode, this command sets the Ethernet IPv4 address, subnet mask, and Gateway on a switch or a chassis.

Command usage depends on the type of IP address and on the platform on which the command is run. Some of the platform- and IP address-specific features of the command are outlined below. For complete details, refer to the *Fabric OS Administrator's Guide*.

Configuring IP Addresses using the command line interface

- The command accepts the **-ipv6** command line syntax with the **--add** or **--delete** option on all platforms that support IPv6 addresses. The **--add** option configures a single static IPv6 address and prefix for the specified managed entity (chassis, CP, or AP). The **--delete** option deletes a static IPv6 address and prefix for the specified managed entity. On modular platforms, the command can be executed only on the active CP.
 - When using the command line syntax to add or delete IPv6 addresses, the managed entity is identified only on modular platforms. To set the CP IPv6 address, use the **-cp** option; to set the IP address for the entire chassis, use the **-chassis** option.
 - When using the command line syntax to add or delete IPv6 addresses on standalone platforms, the implied entity is the single managed entity supported by the platform and must be left unspecified.
 - Additionally, the **-eth0**, **-eth1**, and **-gate** command line options are available with the **--add** or **--delete** option on platforms with blade processors to set the BP Ethernet or Gateway addresses. On a chassis with a blade processor the values for the blade in can be set from the command line using the **-slot** option. The **-slot** option is not accepted in standalone application processors with a hidden blade, such as the AP7600.
- Use the **-auto** and **-noauto** options to enable or disable stateless IPv6 autoconfiguration.
- Use the **-ls** option with appropriate arguments to set or delete the IPv4 Fibre Channel address and prefix for the IPFC interface of a logical switch. In a Virtual Fabric environment, each logical fabric is represented by a separate IPFC network interface. Each of these network interfaces can be assigned a unique IPv4 FC address and prefix. The logical switches that make up a logical fabric are identified by the fabric ID (FID) that is assigned to each of the logical switch instances.

When setting the IPFC interface of a switch that is not in Virtual Fabric mode, use the **-ls** option with FID 128. FID 128 identifies the switch when Virtual Fabrics are disabled.

Setting IP addresses interactively (IPv4 Ethernet address only):

- To set the CP Ethernet IPv4 address, use the **-cp** option; to set the Ethernet IP address for the entire chassis, use the **-chassis** option. When setting the chassis IP address, the command prompts for the Ethernet IP address and Ethernet subnet mask. When setting the CP Ethernet IP address, the command prompts for the host name, Ethernet IP address, Ethernet subnet mask, and Gateway IP address. Valid switch and CP numbers depend on the platform on which the command is run. The command must be executed on the active CP.
- On most standalone platforms (with the exception of the AP7600), **ipAddrSet** runs interactively if invoked without operands. The command prompts for the Ethernet IP address, Ethernet subnet mask, and Gateway IP address. In addition, the command prompts for a specification of whether the Dynamic Host Control Protocol (DHCP) should be used to acquire the Ethernet IP address, Ethernet subnet mask and Gateway IP address. Valid entries are "On" to enable DHCP and "Off" to disable it. When DHCP is enabled, any user-configured Ethernet IP address, Ethernet subnet mask or Gateway IP address is ignored.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

If the IP address is changed while there is an Ethernet connection, connectivity may be lost. The reason for this is explained in detail in the *Fabric OS Administrator's Guide*. If this happens, take one of the following actions to recover the interface:

- Unplug the network cable, wait 5 seconds, and then plug it back in.
- Perform a High Availability (HA) failover routine or hareboot for fixed port switches.
- Power down the switch and then power it back up again.
- Re-establish new Telnet or SSH sessions for the disconnected session after IP address change.

Operands When used in command line mode, the following operands are supported:

- ipv6** Specifies IP address type as static IPv6 including prefix as needed.
- add x:x:x:x:x/n** Sets the specified IPv6 address.
- delete [x:x:x:x:x/n]** Deletes the specified IPv6 address. If no address is specified, any existing IPv6 addresses and prefixes are deleted from the specified or implied entity.
- cp cp_number** Specifies the CP on a chassis. Valid values are **0** or **1**.
- chassis** Specifies the IPv6 address and prefix of a chassis.

On platforms with blade processors, the following additional command line options are supported with the **--add** or **--delete** options:
- eth0 | -eth1** Specifies the local IPv4 address of the blade processor. A prefix is required.
- gate** Specifies the IPv4 address of the blade processor (BP) Gateway (no prefix).

-slot *number*

On a chassis with a blade processor (BP), specifies the slot number. On standalone platforms with a hidden BP, such as the AP76500, this parameter is not accepted.

-ipv6 [-auto | -noauto]

Enables or disables stateless IPv6 autoconfiguration on a switch or chassis. When autoconfiguration is enabled, the host automatically performs configuration of IPv6 addresses and periodic nondisruptive reconfiguration. By default, autoconfiguration is disabled.

-ls *FID*

Specifies the logical fabric ID for which to configure an IPFC network interface. The FID is a decimal number. A switch that is not in Virtual Fabric mode uses the **-ls** parameter with *FID* 128 (the effective, single Logical Fabric number on such switches) to set the IPv4 FC address. Note that setting the IP address for the logical switch is only for the IPFC interface, not for the Ethernet interface.

--add

Assigns a specified IPv4 FC address and prefix to the logical switch instance represented by the specified FID. This command replaces any existing FC IPv4 address.

IPv4_address/prefix

Specifies the IPv4 address and prefix for the IPFC network interface. The IP Address is represented by a dotted decimal number, followed by a slash and a prefix. This operand is required with the **--add** option.

--delete

Deletes the IPv4 FC address and prefix from a logical switch.

-ls *IFID*

Specifies the fabric ID that identifies the logical switch for which to delete the FC IPv4 address and prefix. This operand is required when deleting an FC IPv4 address from a logical switch. On a switch that is not in Virtual Fabric mode, use the **-ls** parameter with *FID* 128 (the effective, single Logical Fabric number on such switches) to delete the IPv4 FC address.

When used interactively to configure IPv4 addresses on a modular platform, ipAddrSet prompts for the following parameters:

-cp *cp_number*

Specifies the managed entity as a CP. Valid values include the following:

0

Sets the Ethernet IP address, Ethernet subnet mask, gateway IP address and host name of CP0.

1

Sets the Ethernet IP address, Ethernet subnet mask, gateway IP address and host name of CP1.

-chassis

Specifies the managed entity as the chassis.

| | |
|---------------------------------|---|
| -ipv4 | Specifies IP address type as static IPv4. |
| -add | Sets the specified IPv4 address. |
| -host <i>hostname</i> | Sets the hostname. |
| -ethip <i>ipaddress</i> | Sets the ethernet IP address. |
| -ethmask <i>mask</i> | Sets the ethernet mask. |
| -gwyip <i>gateway_ip</i> | Sets the gateway IP address. |
| -dhcp[ON OFF] | Enables or disables DHCP. |
| -dhcpv6 | Enables DHCPv6. DHCPv6 is supported only on Brocade 6547 and Brocade 6549. |
| -nodhcpv6 | Disables DHCPv6. DHCPv6 is supported only on Brocade 6547 and Brocade 6549. |
| --clear | Clears all IP (IPv4 and IPv6) addresses on a management interface. |
| --help | Displays command usage. |

Examples

To configure an IPv6 address and prefix on a standalone platform:

```
switch:admin> ipaddrset -ipv6 --add \
fec0:60:69bc:60:260:69ff:fed0:107/64
```

To configure an IPv6 address and prefix on a single CP of a chassis:

```
switch:admin> ipaddrset -cp 0 -ipv6 --add \
1080::8:800:200C:417A/64
```

To delete any existing IPv6 address and prefix on CPO on an enterprise-class platform:

```
switch:admin> ipaddrset cp 0 -ipv6 --delete
```

To configure an IPv4 FC address for the IPFC interface associated with a logical switch with fabric ID 123:

```
switch:admin> ipaddrset -ls 123 --add 11.1.2.4/24
IP address is being changed...Done.
```

To verify the IPv4 FC address for the logical switch:

```
switch:admin> ipaddrshow

CHASSIS
Ethernet IP Address: 10.32.220.10
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.32.220.11
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.32.208.1

CP1
Ethernet IP Address: 10.32.220.12
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.32.208.1
IPFC address for logical fabric ID 128: 1.2.3.4/24

Backplane IP address of CP0 : 10.0.0.5
Backplane IP address of CP1 : 10.0.0.6
```

To delete the IPv4 address for the IPFC interface associated with a logical switch with Virtual Fabric ID 67:

```
switch:admin> ipaddrset -ls 67 --delete
IP address is being changed...Done.
```

To configure an IPv4 FC address for the IPFC interface associated with a switch that is not in Virtual Fabric mode:

```
switch:admin> ipaddrset -ls 128 --add 10.32.72.70/24
IP address is being changed...Done.
```

To verify the changes:

```
switch:admin> ipaddrshow

SWITCH
Ethernet IP Address: 10.32.72.9
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.32.64.1
DHCP: Off
IPFC address for virtual fabric ID 128: 10.32.72.70/24
```

To set the IPv4 address details for a switch chassis in interactive mode:

```
switch:admin> ipaddrset -chassis
Ethernet IP Address [192.168.166.148]:
Ethernet Subnetmask [255.255.255.0]:
Committing configuration...Done.
```

To enable DHCP on a standalone, non-AP platform:

```
switch:admin> ipaddrset
Ethernet IP Address [192.168.74.102]:
Ethernet Subnetmask [255.255.255.0]:
Gateway IP Address [192.168.74.1]:
DHCP [Off]: on
```

To enable DHCPv6 on a device:

```
switch:admin> ipaddrset -ipv6 -dhcpv6
```

To clear all IP address information on a management interface and verify the changes:

```
switch:admin> ipaddrset --clear
Warning: All management interface will be closed, \
which will result in all network connections \
being terminated.
Please initiate a connection to the serial console of \
the switch to configure the management interface.
Would you like to continue with clearing the IP \
configuration?(y/n)y
switch:admin> ipaddrshow
CHASSIS
Ethernet IP Address: none
Ethernet Subnetmask: none

CP0
Ethernet IP Address: none
Ethernet Subnetmask: none
Host Name: cp0
Gateway IP Address: none

CP1
Ethernet IP Address: none
Ethernet Subnetmask: none
Host Name: cp1
Gateway IP Address: none
```

SEE ALSO [ipAddrShow](#)

ipAddrShow

Displays IP address information for a switch or control processor (CP).

Synopsis `ipaddrshow`

`ipaddrshow [-cp cp_number] | -chassis`

`ipaddrshow -slot slot [-eth0 | -eth1 | -gate]`

Description

Use this command to display the IP addresses configured in the system.

The `-cp` option displays the IP address for a specified CP on modular platforms, or use the command without arguments to display the IP address on a standalone switch, or the IP addresses for both CPs on a chassis.

On a standalone switch, the command displays the following information:

- Ethernet IP Address
- Ethernet Subnet mask
- The Gateway IP Address
- Dynamic Host Control Protocol (DHCP): on or Off
- IPv6 Autoconfiguration Enabled: Yes or No
- Local IPv6 Addresses
- IPv6 Gateway address
- DHCPv6: on or off

On modular platforms, the command displays the following information:

For the chassis:

- Ethernet IP Address
- Ethernet Subnet mask

For each CP:

- Ethernet IP Address
- Ethernet Subnet mask
- Host Name
- Gateway IP Address

If the IPFC network interface is configured for logical switches:

- IPFC address for Virtual Fabric ID

For each CP:

- Backplane IP address
- IPv6 Autoconfiguration Enabled: Yes or No. If enabled, the command displays:
 - All local IPv6 Addresses
 - Gateway IP addresses for both CPs

Local IPv6 addresses display the following identifiers:

- IP Address type:
 - **static** - A statically configured IPv6 address.
 - **stateless** - Acquired through stateless autoconfiguration.
- IP Address state:
 - tentative
 - preferred
 - deprecated

Refer to the RFC 2462 specification for more information.

On modular platforms with intelligent blades, the addresses configured for each slot can be shown with the **-slot** option. On standalone platforms, all command options are ignored.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

- cp *cp_number*** On dual-CP systems, specifies the CP card number to be displayed (0 or 1).
- chassis** On dual-CP systems, displays the IP addresses for the chassis.
- slot *slot*** Specifies the slot for a blade.
- eth0 | -eth1 | -gate** For a specified slot, shows only the selected Ethernet interface or the gateway . This operand is optional with the **-slot** option.

Examples To display the IP addresses on a standalone switch:

```
switch:admin> ipaddrshow

SWITCH
Ethernet IP Address: 10.20.21.204
Ethernet Subnetmask: 255.255.240.0
Gateway IP Address: 10.20.16.1
DHCP: Off
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
dhcpv6 fd00:60:69bc:82:a6f3:568b:a9cc:a149/64 preferred
stateless fd00:60:69bc:82:205:33ff:fe94:771f/64 preferred
IPv6 Gateways:
fe80::21b:edff:fe0b:2400
fe80::21b:edff:fe0b:7800
DHCPv6: On
```

To display all IP addresses on a Brocade DCX backbone:

```
switch:admin> ipaddrshow

Chassis
Ethernet IP Address: 10.33.60.85
Ethernet Subnetmask: 255.255.240.0

CP0
Ethernet IP Address: 10.33.60.86
Ethernet Subnetmask: 255.255.240.0
Host Name: cp0
Gateway IP Address: 10.33.48.1

CP1
Ethernet IP Address: 10.33.60.87
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.33.48.1

Backplane IP address of CP0 : 10.0.0.5
Backplane IP address of CP1 : 10.0.0.6
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
chassis 0 stateless fd00:60:69bc:63:205:1eff:fe39:e45a/64 preferred
chassis 0 stateless fec0:60:69bc:63:205:1eff:fe39:e45a/64 preferred
cp 0 stateless fd00:60:69bc:63:205:1eff:fe40:6230/64 preferred
cp 0 stateless fec0:60:69bc:63:205:1eff:fe40:6230/64 preferred
cp 1 stateless fd00:60:69bc:63:205:1eff:fe39:ff2a/64 preferred
cp 1 stateless fec0:60:69bc:63:205:1eff:fe39:ff2a/64 preferred
IPv6 Gateways:
cp 0 fe80:60:69bc:63::3
cp 0 fe80:60:69bc:63::1
cp 0 fe80:60:69bc:63::2
cp 1 fe80:60:69bc:63::1
cp 1 fe80:60:69bc:63::2
cp 1 fe80:60:69bc:63::3
```

To display the IP addresses for a chassis:

```
switch:admin> ipaddrshow -chassis
CHASSIS
Ethernet IP Address: 10.32.220.10
Ethernet Subnetmask: 255.255.240.0
IPv6 Autoconfiguration Enabled: Yes
Local IPv6 Addresses:
chassis 0 stateless fd00:60:69bc:63:205:1eff:fe39:e45a/64 preferred
chassis 0 stateless fec0:60:69bc:63:205:1eff:fe39:e45a/64 preferred
IPv6 Gateways:
```


To display only the IP addresses for CP 1:

```
switch:admin> ipaddrshow

CP1
Ethernet IP Address: 10.33.60.87
Ethernet Subnetmask: 255.255.240.0
Host Name: cp1
Gateway IP Address: 10.33.48.1
```

SEE ALSO [ipAddrSet](#)

ipFilter

Manages the IP filter policies.

Synopsis

```
ipfilter --create polycyname -type ipv4 | ipv6
ipfilter --clone polycyname -from src_polycyname
ipfilter --show [-a] [polycyname]
ipfilter --save [polycyname]
ipfilter --activate polycyname
ipfilter --delete polycyname

ipfilter --addrule polycyname -rule rule_number [-sip source_IP]
    -dp destination_port -proto protocol -act permit | deny
    [-type INPUT | FWD] [-dip destination_IP]
ipfilter --delrule polycyname -rule rule number
ipfilter --transabort
ipfilter --clrcounters
ipfilter --showcounters
```

Description Use this command to manage IP filter policies. The **ipFilter** command and command options are noninteractive, except when prompting for a confirmation.

The IP filter policy sets up a packet filtering firewall to provide access control on the management IP interface. The IPv4 and IPv6 policies are either in the defined configuration or in the active configuration.

Excluding the default policies, there can be a maximum of six policies in the defined configuration and one policy per IPv4 and IPv6 type in the active configuration.

The active policy must be the default policy or one of the policies in the defined configuration. Only the active policies are enforced. All of the **ipFilter** options except **--show** and **--transabort**, create a transaction owned by the management session initiating the commands.

An open transaction prevents other transactions from being created on different management sessions. The **--create**, **--clone**, **--delete**, **--addrule**, and **--delrule** operands modify policies in memory buffer, while operands, **--save** and **--activate** commit policies to the persistent configuration. The operands, **--save** and **--activate**, implicitly end the transaction if all policy changes are committed. The operand **--transabort** explicitly ends an open transaction and aborts policy changes in memory buffer. Closing the management session that owns the transaction also aborts policy changes and closes the transaction.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

In a Virtual Fabric environment, IP Filter policies are treated as chassis-wide configurations and apply to all logical switches in the chassis. Chassis permissions are required to manage IP Filter policies.

Operands This command has the following operands:

policyname

Specifies an IP filter policy name. The policy name is a unique string composed of a maximum of 20 alphanumeric or underscore characters. The `default_ipv4` and `default_ipv6` names are reserved for default IP filter policies. The policy name is case-insensitive and is always stored as lower case. The policy type identifies the policy as an IPv4 or IPv6 filter. You can create a maximum of eight IP filter policies.

`--create policyname -type ipv4 | ipv6`

Creates an IP filter policy with the specified name and type. The policy created is stored in a temporary buffer and is lost if the policy is not saved to the persistent configuration.

`--clone policyname -from src_policyname`

Creates a replica of an existing IP filter policy. The cloned policy is stored in a temporary buffer and has the same rules as the original policy.

`--show [-a] [policyname]`

Displays the IP filter policy content for the specified policy name or all IP filter policies if *policyname* is not specified. For each IP filter policy, the policy name, type, persistent state, and policy rules are displayed. The policy rules are listed by the rule number in ascending order. The `-a` option displays the traffic type and the configured destination IP.

Command output displays without pagination. Use `command | more` to display the output with page breaks. If a temporary buffer exists for an IP filter policy, the `--show` operand displays the content in the temporary buffer, with the persistent state set to modified defined or modified active.

`--save [policyname]`

Saves one or all IP filter policies persistently as the defined configuration. This operand is optional. If a policy name is specified, only the specified IP filter policy in the temporary buffer is saved; otherwise, all IP filter policies in the temporary buffer is saved. Only the CLI session that owns the updated temporary buffer can run this command. Modification to an active policy cannot be saved without being applied. Therefore, the `--save` option is blocked for the active policies; instead use the `--activate` option.

`--activate policyname`

Activates the specified IP filter policy. IP filter policies are not enforced until they are activated. Only one IP filter policy per IPv4 and IPv6 type can be active. If there is a temporary buffer for the policy, the policy is saved to the defined configuration and activated at the same time. If there is no temporary buffer for the policy, the policy existing in the defined configuration becomes active. The policy to be activated replaces the existing active policy of the same type. Activating the default IP filter policies returns the IP management interface to its default state. An IP filter policy without any rule cannot be activated. This operand prompts for confirmation before proceeding.

`--delete policyname`

Deletes the specified IP filter policy. Deleting an IP filter policy removes it from the temporary buffer. To permanently delete the policy from the persistent database, issue the `ipfilter --save` command. An active IP filter policy cannot be deleted.

- addrule *polycyname***
Adds a new rule to the specified IP filter policy. The change made to the specified IP filter policy is not saved to the persistent configuration until saved or activated.
- The following arguments are supported with the **--addrule** option:
- sip *source_IP***
Specifies the source IP address. For filters of type IPv4, the address must be a 32-bit address in dot notation, or a CIDR-style IPv4 prefix. For filters of type IPv6, the address must be a 128-bit IPv6 address in any format specified by RFC3513, or a CIDR-style IPv6 prefix. The source IP option is not supported for FORWARD traffic
- dp *destination_port***
Specifies the destination port number, a range of port numbers, or a service name. Note that blocking or permitting of ports 1024 and above is not allowed. These ports are used by various applications and services on the switch.
- proto *protocol***
Specifies the protocol type, for example, *tcp* or *udp*.
- act permit | deny**
Specifies the permit or deny action associated with this rule. Blocking or permitting port 1024 and above is not allowed. Ports numbered 1024 and higher are used by applications for services such as FTP and blocking these ports may cause these applications to behave in unexpected ways. Blocking or permitting the RPC ports 111, 897, and 898 is not allowed.
- rule *rule_number***
Adds a new rule at the specified rule index number. The rule number must be between 1 and the current maximum rule number plus one.
- type INPUT | FWD**
Specifies the type of traffic that is allowed for the specified IP address. Forwarding rules manage the bidirectional traffic between the external Ethernet interface (eth0/bond0) and the inband management interface (inbd+). INPUT traffic is the default type of traffic for IP filter rules.
- dip *destination_IP***
Specifies the destination IP address. For filters of type IPV4, the address must be a 32-bit address in dot notation, or a CIDR-style IPv4 prefix. For filters of type IPv6, the address must be in a 128-bit IPv6 address in any format specified by RFC3513, or a CIDR-style IPv6 prefix. The destination IP option is not supported for INPUT traffic type.
- delrule *polycyname* -rule *rule_number***
Deletes a rule from the specified IP filter policy. Deleting a rule in the specified IP filter policy causes the rules following the deleted rule to shift up in rule order. The change to the specified IP filter policy is not saved to the persistent configuration until it is saved or activated.
- transabort**
A transaction is associated with a CLI or manageability session, which is opened implicitly when you execute the **--create**, **--addrule** and **--delrule** subcommands. The **--transabort** command explicitly ends the transaction owned by the current CLI or manageability session. If a transaction is not ended, other CLI or manageability sessions are blocked on the subcommands that would open a new transaction.

--clrcounters

Clears the IP filter counters. This command requires root permissions.

--showcounters

Displays the IP filter counters. This command requires root permissions.

Examples

To create an IP filter for a policy with an IPv6 address:

```
switch:admin> ipfilter --create ex1 -type ipv6
```

To add a new rule to the policy and specify the source IP address, destination port, and protocol, and to permit the rule:

```
switch:admin> ipfilter --addrule ex1 \
-sip fec0:60:69bc:60:260:69ff:fe80:d4a -dp 23 \
-proto tcp -act permit
```

To display all existing IP filter policies:

```
switch:admin> ipfilter --show
```

```
Name: default_ipv4, Type: ipv4, State: active
Rule  Source IP      Protocol  Dest Port  Action
1     any               tcp       22         permit
2     any               tcp       23         permit
3     any               tcp       80         permit
4     any               tcp       443        permit
5     any               udp       161        permit
6     any               udp       123        permit
7     any               tcp       600 - 1023 permit
8     any               udp       600 - 1023 permit
```

```
Name: default_ipv6, Type: ipv6, State: active
Rule   Source IP      Protocol  Dest Port  Action
1     any             tcp       22         permit
2     any             tcp       23         permit
3     any             tcp       80         permit
4     any             tcp       443        permit
5     any             udp       161        permit
6     any             udp       123        permit
7     any             tcp       600 - 1023 permit
8     any             udp       600 - 1023 permit
```

To activate the IP Filter policy "ex1":

```
switch:admin> ipfilter --activate ex1
```

To display all IP Filter policies, including the activated policy:

```
switch:admin> ipfilter --show
```

```
Name: default_ipv4, Type: ipv4, State: active
Rule   Source IP      Protocol  Dest Port  Action
1     any             tcp       22         permit
2     any             tcp       23         permit
3     any             tcp       80         permit
4     any             tcp       443        permit
5     any             udp       161        permit
```

ipFilter

```
6    any                udp      123      permit
7    any                tcp      600 - 1023 permit
8    any                udp      600 - 1023 permit
```

```
Name: default_ipv6, Type: ipv6, State: defined
Rule  Source IP          Protocol  Dest Port  Action
1    any                tcp      22         permit
2    any                tcp      23         permit
3    any                tcp      80         permit
4    any                tcp      443        permit
5    any                udp      161        permit
6    any                udp      123        permit
7    any                tcp      600 - 1023 permit
8    any                udp      600 - 1023 permit
```

```
Name: ex1, Type: ipv6, State: active
Rule  Source IP          Protocol  Dest Port  Action
1    fec0:60:69bc:60:260:69ff:fe80:d4a tcp 23         permit
```

To create an IPv4-type IP filter policy:

```
switch:admin> ipfilter --create ex2 -type ipv4
```

To add a rule to the created policy "ex2":

```
switch:admin> ipfilter --addrule ex2 -sip 10.32.69.99 \
               -dp 23 -proto tcp -act permit
```

To display the IP filter policies, including the new policy:

```
switch:admin> ipfilter --show

Name: default_ipv4, Type: ipv4, State: active
Name: default_ipv4, Type: ipv4, State: active
Rule  Source IP          Protocol  Dest Port  Action
1    any                tcp      22         permit
2    any                tcp      23         permit
3    any                tcp      80         permit
4    any                tcp      443        permit
5    any                udp      161        permit
6    any                udp      123        permit
7    any                tcp      600 - 1023 permit
8    any                udp      600 - 1023 permit

Name: ex2, Type: ipv4, State: defined (modified)
Rule  Source IP          Protocol  Dest Port  Action
1    10.32.69.99        tcp      23         permit
```

To save the IP Filter policy "ex2" (the status of the policy changes from modified to defined after the policy is saved):

```
switch:admin> ipfilter --save ex2
```

```
switch:admin> ipfilter --show

Name: default_ipv4, Type: ipv4, State: active
Name: default_ipv4, Type: ipv4, State: active
Rule  Source IP          Protocol  Dest Port  Action
1    any                tcp      22         permit
2    any                tcp      23         permit
```

```

3    any          tcp      80      permit
4    any          tcp      443     permit
5    any          udp      161     permit
6    any          udp      123     permit
7    any          tcp      600 - 1023 permit
8    any          udp      600 - 1023 permit

```

Name: ex2, Type: ipv4, State: defined

```

Rule    Source_IP      Protocol  Dest Port Action
1       10.32.69.99      tcp      23      permit

```

To add a rule for traffic of type FORWARD and to display the result:

```

switch:admin> ipfilter --addrule ex2 -rule 1 -sip \
               10.32.69.99 -dp 23 -proto tcp -act permit -type fwd \
               -dip 10.32.69.90

```

```

switch:admin> ipfilter --show -a

```

Legends: [Proto-Protocol, P-Permit, D-Deny, FWD-FORWARD, I/P-INPUT]

Name: default_ipv4, Type: ipv4, State: active

```

Rule    Source_IP  Proto  Dest_Port  Action  Flow  Destination_IP
1       any         tcp    22         P      I/P   any
2       any         tcp    23         P      I/P   any
3       any         tcp    80         P      I/P   any
4       any         tcp    443        P      I/P   any
5       any         udp    161        P      I/P   any
6       any         udp    123        P      I/P   any
7       any         tcp    600-1023  P      I/P   any
8       any         udp    600-1023  P      I/P   any

```

Name: default_ipv6, Type: ipv6, State: active

```

Rule    Source_I  Proto  Dest_Port  Action  Flow  Destination_IP
1       any         tcp    22         P      I/P   any
2       any         tcp    23         P      I/P   any
3       any         tcp    80         P      I/P   any
4       any         tcp    443        P      I/P   any
5       any         udp    161        P      I/P   any
6       any         udp    123        P      I/P   any
7       any         tcp    600-1023  P      I/P   any
8       any         udp    600-1023  P      I/P   any

```

Name: ex2, Type: ipv4, State: defined (modified)

```

Rule    Source_IP  Proto  Dest_Port  Action  Flow  Destination_IP
1       10.32.69.99  tcp    23         P      FWD   10.32.69.99

```

SEE ALSO [distribute](#)

ipSecConfig

Configures Internet Protocol security (IPSec) policies for Ethernet management interfaces.

Synopsis

```
ipsecconfig --enable [default] --disable
ipsecconfig --add | --modify type [subtype] [arguments]
ipsecconfig --delete [type] arguments
ipsecconfig --flush manual-sa
ipsecconfig --show type [subtype] arguments
ipsecconfig --help [command_type subtype]
```

Description Use this command to configure the Internet Protocol Security (IPSec) feature for traffic flows on switch Ethernet management interfaces, or to display the current configuration.

Internet Protocol security (IPSec) is a framework of open standards that provides private, secure communication over Internet Protocol (IP) networks through the use of cryptographic security services.

IPSec uses different protocols to ensure the authentication, integrity, and confidentiality of the communication.

- Encapsulating Security Payload (ESP) provides confidentiality, data integrity and data source authentication of IP packets, and protection against replay attacks.
- Authentication Header (AH) provides data integrity, data source authentication, and protection against replay attacks, but unlike ESP, AH does not provide confidentiality.

IPSec can protect either the entire IP datagram or only the upper-layer protocols. The appropriate modes are called tunnel mode and transport mode.

- In tunnel mode the IP datagram is fully encapsulated by a new IP datagram using the IPSec protocol.
- In transport mode only the payload of the IP datagram is handled by the IPSec protocol; it inserts the IPSec header between the IP header and the upper-layer protocol header.

The IPSec key management supports Internet Key Exchange (IKE) or Manual key/SA entry.

- In IKE the IPSec protocol negotiates shared security parameters and keys. Security Associations (SAs) used in IKE use automatically generated keys for authentication negotiation between peers.
- Manual key/SA entry requires the keys to be generated and managed manually, and it is therefore suited for small static environments. For the selected authentication or encryption algorithms, the correct keys must be generated. The key length is determined by the algorithm selected. Refer to the *Fabric OS Administrator's Guide* for more information.

The following IPSec configuration tasks can be performed with this command:

- Enable or disable the IPSec policies.
- Configure IP address for both IPv4 and IPv6 format.
- Configure three types of policies and their respective components:
 - IPSec policy including selector, transform, SA-proposal, and SA.
 - IKE policy (automatic key management).
 - Manual SA (manual SA management).
- Modify existing IPSec and IKE policies.
- Delete existing policies and SAs from the configuration database.
- Flush existing SAs from the kernel SA database (SADB).
- Display policy parameters.

Representation of IP addresses

When configuring IPSec policies, IP addresses and ports must be specified in the following format:

IP address

IPv4 addresses are expressed in dotted decimal notation consisting of numeric characters (0-9) and periods (.), for example, 203.178.141.194.

IPv6 address consist of hexadecimal digits (09afAF), colons (:) and a percent sign (%) if necessary, for example, 2001:200:0:8002:203:47ff:fea5:3085

network prefix

A network prefix is represented by a number followed by a slash (/), for example, 1/0.

Notes IPSec configuration changes take effect upon execution and are persistent across reboot.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command does not provide IPSec protection for traffic flows on external management interfaces of intelligent blades in a chassis, nor does it support protection of traffic flows on FCIP interfaces.

This command does not support manipulating preshared keys corresponding to the identity of the IKE peer or group of peers. Use **secCertUtil** to import, delete, or display the preshared keys in the local switch database.

The MD5 hash algorithm is blocked when FIPS mode is enabled.

Refer to the Examples section for specific use cases and associated command sequences. Refer to the *Fabric OS Administrator's Guide* for configuration procedures.

This command accepts abbreviated operands. The abbreviated string must contain the minimum number of characters necessary to uniquely identify the operand within the set of available operands.

Operands

This command has the following operands:

--enable

Enables IPsec on the switch. Existing IPsec configurations are enabled by this command. IPsec is disabled by default. It must be enabled before you can configure the policies and parameters. The following operand is optional:

default

Clears the existing policies (automatic key management and manual keyed entries) and resets the configuration databases to default values.

--disable

Disables IPsec on the switch. All active TCP sessions are terminated when you disable IPsec.

--add | --modify

Adds or modifies an IPsec or IKE policy in an existing enabled configuration. Not all parameters can be modified. Parameters that cannot be modified are indicated below. When modifying a policy the names and identifiers need to refer to valid existing entities. The syntax is as follows:

```
--add | --modify type [subtype] [arguments]
```

type

Specifies the policy to be created. Supported policies include the following:

policy ips

Creates or modifies an IPsec policy. This policy determines the security services afforded to a packet and the treatment of a packet in the network. An IPsec policy allows classifying IP packets into different traffic flows and specifies the actions or transformations performed on IP packets on each of the traffic flows. The main components of an IPsec policy are: IP packet filter/selector (IP address, protocol, and port information) and transform set.

subtype

A *subtype* is required when configuring an IPsec policy. The subtype specifies the components to be configured. The following are required subtypes for the IPsec policy:

selector

Creates a selector that is applied to the IP data traffic. A selector consists of a set of parameters that identify the IP traffic that needs IPsec protection. To configure the selector, the following parameters must be specified:

-tag name

Specifies a name for the selector. This is a user-generated name. The name must be between 1 and 32 characters in length, and may include alphanumeric characters, dashes (-), and underscores (_).

-direction in | out

Specifies traffic flow direction as inbound or outbound.

-local IP_address[prefixlength]

Specifies the source IPv4 or IPv6 address.

-remote *IP_address*[/*prefixlength*]
Specifies the peer IPv4 or IPv6 address

-transform *name*
Specifies the transform to be included in the selector. You must create the transform before you can use in the selector. Use **ipsecConfig --show policy ips transform** to display existing transforms.

-protocol *protocol_name*
Specifies the upper layer protocols to be selected for IPSec protection. Valid protocols include **tcp**, **udp**, **icmp** or **any**. When any is specified all existing protocols are selected for protection. This operand is optional.

transform

Creates the IPSec transform set. The transform set is a combination of IPSec protocols and cryptographic algorithms that are applied on the packet after it is matched to a selector. The transform set specifies the IPSec protocol, the IPSec mode, and the action to be performed on the IP packet. It also specifies the key management policy that is needed for the IPSec connection and the encryption and authentication algorithms to be used in security associations when IKE is used as key management protocol. The following operands are required:

-tag *name*
Specifies a name for the transform. This is a user-generated name. The name must be between 1 and 32 characters in length, and may include alphanumeric characters, dashes (-), and underscores (_).

-mode tunnel | transport
Specifies the IPSec transform mode. In tunnel mode, the IP datagram is fully encapsulated by a new IP datagram using the IPSec protocol. In transport mode, only the payload of the IP datagram is handled by the IPSec protocol inserting the IPSec header between the IP header and the upper-layer protocol header.

-sa-proposal *name*
Specifies the SA proposal to be included in the transform. You must create the SA proposal first before you can include it in the transform. Use **ipsecConfig --show policy ips sa-proposal -a** for a listing of existing SA proposals.

-action discard | bypass | protect
Specifies the protective action the transform should take regarding the traffic flows.

-ike *name*
Specifies the IKE policy to be included in the transform. This operand is optional. Use **ipsecConfig --show policy ike -a** for a listing of existing IKE policies.

-local *IP_address*[/*prefixlength*]
Specifies the source IPv4 or IPv6 address. This operand is optional. If a local source IP address is defined, a remote peer IP address must also be defined.

-remote *IP_address*[/*prefixlength*]
Specifies the peer IPv4 or IPv6 address. This operand is optional. If a remote peer IP address is defined, a local source IP address must also be defined.

sa-proposal

Defines the security associations (SA) proposal, including name, SAs to be included and lifetime of the proposal. The following operands are supported:

-tag *name*

Specifies a name for the SA proposal. This is a user-generated name. The name must be between 1 and 32 characters in length, and may include alphanumeric characters, dashes (-), and underscores (_).

-sa *name[,name]*

Specifies the SAs to include in the SA proposal. The bundle consists of one or two SA names, separated by commas. For SA bundles, [AH, ESP] is the supported combination. The SAs must be created prior to being included in the SA proposal. This operand is required.

-lttime *number*

Specifies the SA proposal's lifetime in seconds. This operand is optional. If a lifetime is not specified, the SA does not expire. If lifetime is specified both in seconds and in bytes, the SA expires when the first expiration criterion is met.

-ltbyte *number*

Specifies the SA proposal's lifetime in bytes. The SA expires after the specified number of bytes have been transmitted. This operand is optional.

sa

Defines the Security Association. An SA specifies the IPsec protocol (AH or ESP), the algorithms used for encryption and authentication, and the expiration definitions used in security associations of the traffic. IKE uses these values in negotiations to create IPsec SAs.

You cannot modify an SA once it is created. Use **ipsecConfig --flush manual-sa** to remove all SA entries from the kernel SA database (SADB) and start over.

-tag *name*

Specifies a name for the SA. This is a user-generated name. The name must be between 1 and 32 characters in length, and may include alphanumeric characters, dashes (-), and underscores (_). This operand is required.

-protocol *ah | esp*

Specifies the IPsec protocol. Encapsulating Security Payload (ESP) provides confidentiality, data integrity and data source authentication of IP packets, and protection against replay attacks. Authentication Header (AH) provides data integrity, data source authentication, and protection against replay attacks but, unlike ESP, does not provide confidentiality. This operand is required.

-auth *algorithm*

Specifies the authentication algorithm. This operand is required. Valid algorithms include the following:

hmac_md5

MD5 authentication algorithm

hmac_sha1

SHA1 authentication algorithm

-enc *algorithm*

Specifies the encryption algorithm. This operand is required. Valid algorithms include the following:

3des_cbc

3DES encryption algorithm

blowfish_cbc

Blowfish encryption algorithm

null_enc

Null encryption algorithm

aes256_cb

AES-256 algorithm

-spi *number*

Specifies the security parameter index (SPI) for the SA. This is a user-defined index. Valid SPI numbers consist of numeric characters (0-9). This operand is optional.

policy ike

Creates or modifies an IKE policy configuration. No *subtype* is required with this command. The command defines the following IKE policy parameters: IKE version, IP address of the remote entity, IP address of the local entity, encryption algorithm, hash algorithm, PRF algorithm, DH group, authentication method, path and filename of the preshared key. The syntax is as follows: **ipsecConfig --add | --modify ike *arguments***.

arguments

Valid arguments for **policy ike** include the following:

-tag *name*

Specifies a name for the IKE policy. This is a user-generated name. The name must be between 1 and 32 characters in length, and may include alphanumeric characters, dashes (-), and underscores (_). This operand is required.

remote *IP_address*[/*prefixlength*]

Specifies the peer IPv4 or IPv6 address and prefix.

-id *identifier*

Specifies the local identifier. The switch is identified by its IPv4 or IPv6 address.

-remoteid *identifier*

Specifies the peer identifier. The remote peer is identified by its IPv4 or IPv6 address.

-enc *algorithm*

Specifies the encryption algorithm. Valid encryption algorithms include the following:

3des_cbc

DES algorithm

blowfish_cbc

Blowfish algorithm

aes128_cbc

AES 128-bit algorithm

aes256_cbc
AES 256-bit algorithm

-hash *algorithm*
Specifies the hash algorithm. Valid hash algorithms include the following:

hmac_md5
MD5 algorithm

hmac_sha1
SH1 algorithm

-prf *algorithm*
Specifies the PRF algorithm. Valid PRF algorithms include the following:

hmac_md5
MD5 algorithm

hmac_sha1
SH1 algorithm

-auth *psk* | *dss* | *rsasig*
Specifies the authentication method as one of the following:

psk
Authenticate using preshared keys.

dss
Authenticate using digital signature standard.

rsasig
Authenticate using an RSA signature.

-dh *number*
Specifies the DH group number as one of the following:

1
Specifies DH group **modp768**.

2
Specifies DH group **modp1024**.

14
Specifies DH group **modp2048**.

The following operands are optional (use **secCertUtil** to import the key files to the local and remote systems)

-psk *file*
Specifies the preshared key filename.

-pubkey *file*
Specifies the public key filename (in X.509 PEM format).

-privkey *file*
Specifies the private key filename (in X.509 PEM format).

-peerpubkey *file*
Specifies the peer public key filename (in X.509 PEM format).

-lttime *number*
Specifies the key lifetime in seconds. If a lifetime is not specified, the keys do not expire. If a lifetime is specified both in seconds and in bytes, the keys expire when the first expiration criterion is met.

-ltbyte *number*
Specifies the key lifetime in bytes. The keys expire after the specified number of bytes have been transmitted.

-pfs on | off
Enables or disables Perfect Forward Secrecy (PFS). PFS is disabled by default. When PFS is disabled, IKE uses the initial master key it generates in Phase1 to generate the keys for SA connections in Phase2. When PFS is enabled, a new key is generated for keying the SAs. Enabling PFS may provide enhanced protection against keys compromise.

-version 1 | 2
Specifies the IKE version. This operand is optional. If not specified, IKEv2 is used (2). If 1 is specified, IKEv1 is selected. Use **-v 2** to revert to version 2 after version 1 was set.

manual-sa
Creates manually keyed SADB entries. When using this option, you must generate the keys manually. The lifetime of an SA entry created using this command is infinite. You cannot modify manually keyed SA entries. Use **ipsecConfig --flush**, or **ipsecConfig --delete** and recreate the entries. The syntax for creating an SADB entry is as follows:

ipsecconfig --add manual-sa *arguments*.

arguments
Valid arguments for **manual-sa** include the following:

-sp *number*
Specifies the security parameter index (SPI) for the SA. This is a user-defined index. Valid SPI numbers consist of numeric characters (0-9).

-local *ipaddress*
Specifies the local IPv4 or IPv6 address.

-remote *ipaddress*
Specifies the remote IPv4 or IPv6 address.

-protocol *protocol_name*
Specifies the upper layer protocols to be selected for protection. Valid protocols include **tcp**, **udp**, **icmp** or **any**. When any is specified all existing protocols are selected for protection.

-ipsec ah | esp
Specifies the IPSec protocol. Encapsulating Security Payload (ESP) provides confidentiality, data integrity and data source authentication of IP packets, and protection against replay attacks. Authentication Header (AH) provides data integrity, data source authentication, and protection against replay attacks but, unlike ESP, does not provide confidentiality.

-action discard | bypass | protect
Specifies the IPSec protection type regarding the traffic flows.

- direction in | out**
Specifies traffic flow direction as inbound or outbound.
- mode tunnel | transport**
Specifies the IPsec mode. In tunnel mode, the IP datagram is fully encapsulated by a new IP datagram using the IPsec protocol. In transport mode, only the payload of the IP datagram is handled by the IPsec protocol; it inserts the IPsec header between the IP header and the upper-layer protocol header.
- enc algorithm**
Specifies the encryption algorithm. Valid encryption algorithms include the following:
- 3des_cbc**
3DES algorithm
 - null_enc**
Null encryption algorithm (cleartext)
- auth algorithm**
Specifies the authentication algorithm. Valid authentication algorithms include the following:
- hmac_md5**
MD5 algorithm
 - hmac_sha1**
SHA1 algorithm
- enc-key number**
Specifies the encryption key. This is a user-generated key based on the length of the key. Use the LINUX random key generator or any other comparable third party utility to generate the manual SA keys. Refer to the *Fabric OS Administrator's Guide* for details.
- A 192-bit value for the **3des_cbc** encryption algorithm, for example, 0x96358c90783bbfa3d7b196ceabe0536b
 - A zero-bit value for the **null_enc** encryption algorithm.
- auth-key number**
Specifies the authentication key. This is a user-generated key based on the length of the key. Valid keys include the following:
- A 128-bit value for the **hmac_md5** authentication algorithm.
 - A 160-bit value for the **hmac_sha1** authentication algorithm.
- The following operands are optional:
- tunnel-local *ipaddress***
Specifies the local tunnel IPv4 or IPv6 address.
- tunnel-remote *ipaddress***
Specifies the peer tunnel IPv4 or IPv6 address.

--delete

Deletes a specified policy or all policies of a certain *type* from the configuration database. You can delete IPsec policies, IKE policies, and SADB entries. When deleting IPsec policies, you have the option to delete specific components only, such as the transform or the selector, and recreate these components without having to recreate the entire policy. The syntax for deleting a policy is as follows:

ipseconfig --delete *type* [*subtype*] *arguments*

type

Specifies the policy to be deleted. Valid policy types include the following:

policy ips

Deletes a specified IPsec policy or all IPsec policies.

subtype

Optionally specify a component (*subtype*) to delete the component only:

selector

Deletes the selector for a specified IPsec policy, or all selectors of all configured IPsec policies.

transform

Deletes the transform for a specified IPsec policy, or all transforms of all configured IPsec policies.

sa-proposal

Deletes the SA proposal for a specified IPsec policy, or all SA proposals of all configured IPsec policies.

sa

Deletes the SAs for a specified IPsec policy, or all SAs of all configured IPsec policies.

policy ike

Deletes a specified IKE policy or all configured IKE policies.

sa

Deletes a specified SADB entry or all manual SADB entries.

manual-sa

Deletes the SA policy entries used in manually keyed SA entries from the configuration database.

arguments

Specifies the selection as one of the following:

-a

Deletes all configuration information for the specified *type* and *subtype*.

-tag *name*

Deletes all configuration information for the specified policy *type*

| | |
|--------------------------|--|
| --flush manual-sa | Flushes all SA entries (including manually keyed and automatically keyed SAs) from the kernel SADB. All active TCP sessions that are using IPsec protection are terminated when this command is executed. This command, unlike delete, does not remove the policies from the configuration database. Flushing any other policy parameters is not supported. |
| --show | Displays current IPsec or IKE configuration. The syntax for the display option is as follows: --show <i>type</i> [<i>subtype</i>] <i>arguments</i> |
| <i>type</i> | Specifies the policy to be displayed. Valid values for <i>type</i> include the following: |
| policy ips | Displays the IPsec policy configuration. A policy <i>subtype</i> must be specified when displaying the IPsec policy configuration. Valid subtypes include the following: |
| selector | Displays IPsec selector parameters including IPsec policy name, IP address of the local entity, IP address of the remote entity, direction of traffic flow (inbound or outbound), upper layer protocol used, and IPsec transform index. |
| transform | Displays IPsec transform parameters including IPsec policy name, key management protocol (version) or manual SA, processing option for selected IP traffic, IPsec mode (tunnel or transport), IP address of the local entity, IP address of the remote entity, and SA proposal. |
| sa-proposal | Displays the parameters of the SA proposal, including proposal name, lifetime (in seconds and in byte units, or infinite), and associated SA definitions. |
| sa | Displays security association (SA) parameters for the specified IPsec policies including policy names, IPsec protocol used (AH or ESP), encryption and authentication algorithms. |
| policy ike | Displays the IKE policy configuration. No <i>subtype</i> is required with this command. The command displays the following IKE policy parameters: IKE version, IP address of the remote entity, IP address of the local entity, encryption algorithm, hash algorithm, PRF algorithm, DH group, authentication method, path and filename of the preshared key. |
| manual-sa | Displays the Security Associations in the local SADB. No <i>subtype</i> is required with this command. |
| <i>arguments</i> | Specifies the display selection as one of the following: |
| -a | Displays all configuration information for the specified <i>type</i> and <i>subtype</i> . |

-tag *name*

Displays configuration information for the specified IPsec policy only.

--help

Displays the command usage. Optionally use **--help** with a specified command, type, and subtype to display the syntax for specific commands. Do not include a double dash when specifying the command.

Examples **Example 1**

The following example illustrates how to secure traffic between two systems using AH protection with MD5 and configure IKE with preshared keys. The two systems are a switch, BROCADE300 (IPv4 address 10.33.74.13), and an external UNIX server (10.33.69.132).

1. On the system console, log in to the switch as Admin and enable IPsec.

```
switch:admin> ipsecconfig --enable
```

2. Create an IPsec SA policy named AH01, which uses AH protection with MD5.

```
switch:admin> ipsecconfig --add policy ips sa \
-t AH01 -p ah -auth hmac_md5
```

3. Create an IPsec proposal IPSEC-AH to use AH01 as SA.

```
switch:admin> ipsecconfig --add policy ips sa-proposal \
-t IPSEC-AH -sa AH01
```

4. Configure the SA proposal's lifetime in time units.

```
switch:admin> ipsecconfig --add policy ips sa-proposal \
-t IPSEC-AH -lifetime 280000 -sa AH01
```

5. Import the preshared key file (e.g., ipseckey.psk) using the **seCcertUtil import** command.

6. Configure an IKE policy for the remote peer.

```
switch:admin> ipsecconfig --add policy ike -t IKE01 \
-remote 10.33.69.132 -id 10.33.74.13 \
-remoteid 10.33.69.132 -enc 3des_cbc \
-hash hmac_md5 -prf hmac_md5 -auth psk \
-dh modp1024 -psk ipseckey.psk
```

7. Create an IPsec transform named TRANSFORM01 to use transport mode to protect traffic identified for IPsec protection and use IKE01 as the key management policy.

```
switch:admin> ipsecconfig --add policy ips transform \
-t TRANSFORM01 -mode transport \
-sa-proposal IPSEC-AH -action protect -ike IKE01
```

8. Create traffic selectors to select the outbound and inbound traffic that needs to be protected.

```
switch:admin> ipsecconfig --add policy ips selector \
-t SELECTOR-OUT -d out -l 10.33.74.13 -r 10.33.69.132 \
-transform TRANSFORM01
```

```
switch:admin> ipsecconfig --add policy ips selector \
-t SELECTOR-IN -d in -l 10.33.69.132 -r 10.33.74.13 \
-transform TRANSFORM01
```

9. Verify the IPsec SAs created using IKE for above traffic flow using **ipsecConfig --show manual-sa -a**. Refer to the "IPsec display commands" section for an example.

10. Perform the equivalent steps on the remote peer to complete the IPsec configuration. Refer to your server administration guide for instructions.

Example 2

The following example illustrates how to secure traffic between two systems using ESP protection with 3DES_CBC encryption and SHA1 authentication, and how to configure IKE with RSA Certificates signed by the certification authority (CA). The two systems are A SWITCH, BROCADE300 (IPv6 address fe80::220:1aff:fe34:2e82), and an external UNIX host (IPv6 address fe80::205:1fff:fe51:f09e).

1. On the system console, log in to the switch as Admin and enable IPsec.

```
switch:admin> ipsecconfig --enable
```

2. Create an IPsec SA policy named ESP01, which uses ESP protection with 3DES and SHA1.

```
switch:admin> ipsecconfig --add policy ips sa -t ESP01 \
-p esp -enc 3des_cbc -auth hmac_sha1
```

3. Create an IPsec proposal IPSEC-ESP to use ESP01 as the SA.

```
switch:admin> ipsecconfig --add policy ips sa-proposal \
-t IPSEC-ESP -sa ESP01
```

4. Configure the SA proposal lifetime in seconds.

```
switch:admin> ipsecconfig --add policy ips sa-proposal \
-t IPSEC-ESP -lifetime 280000 -sa ESP01
```

5. Import the public key for the BROCADE300 (Brocade300.pem), the private key for BROCADE300 (Brocade300-key.pem), and the public key of the external host (remote-peer.pem) in X.509 PEM format from the remote certificate server (10.6.103.139).

```
switch:admin> seccertutil import -ipaddr 10.103.6.139 \
-remotedir /root/certs -certname Brocade300.pem
```

```
switch:admin> seccertutil import -ipaddr 10.103.6.139 \
-remotedir /root/certs -certname Brocade300-key.pem
```

```
switch:admin> seccertutil import -ipaddr 10.103.6.139 \  
-remotedir /root/certs -certname remote-peer.pem
```

6. Import the CA certificate that was used to sign the public certificates of BROCADE300 and the remote peer as IPSECCA.pem.

```
switch:admin> seccertutil import -ipaddr 10.103.6.139 \  
-remotedir /root/certs -certname IPSECCA.pem
```

7. Configure an IKE policy for the remote peer UNIX host.

```
switch:admin> ipsecconfig --add policy ike -t IKE01 -remote \  
fe80::205:1fff:fe51:f09e -id fe80::220:1aff:fe34:2e82 \  
-remoteid fe80::205:1fff:fe51:f09e \  
-enc 3des_cbc -hash hmac_md5 -prf hmac_md5 -auth rsasig \  
-dh modp1024 -pubkey "Brocade300.pem" \  
-privkey "Brocade300-key.pem" -peerpubkey "remote-peer.pem"
```

8. Create an IPSec transform TRANSFORM01 to use transport mode to protect traffic identified for IPSec protection and use IKE01 as the key management policy.

```
switch:admin> ipsecconfig --add policy ips transform \  
-t TRANSFORM01 -mode transport -sa-proposal IPSEC-ESP \  
-action protect -ike IKE01
```

9. Create traffic selectors to select outbound and inbound TCP traffic that needs to be protected.

```
switch:admin> ipsecconfig --add policy ips selector \  
-t SELECTOR-OUT -d out -l fe80::220:1aff:fe34:2e82 \  
-r fe80::205:1fff:fe51:f09e \  
-protocol "tcp" -transform TRANSFORM01
```

```
switch:admin> ipsecconfig --add policy ips selector \  
-t SELECTOR-IN -d in -l fe80::205:1fff:fe51:f09e -r \  
fe80::220:1aff:fe34:2e82 -protocol "tcp" -t transform TRANSFORM01
```

10. Verify the IPSec SAs using `ipSecConfig --show manual-sa -a`. Refer to the "IPSec display commands" section for an example.
11. Perform the equivalent steps on the remote peer to complete the IPSec configuration. Refer to your server administration guide for instructions.

Example 3

The following example illustrates how to secure traffic between two systems using AH with SHA1 and ESP protection with 3DES and configure IKE with preshared keys. The two systems are a switch, a Brocade 300 (IP address 10.33.74.13), and an external UNIX host (IPv4 address 10.33.69.132).

1. On the system console, log in to the switch as Admin and enable IPsec.

```
switch:admin> ipsecconfig --enable
```

2. Create an IPsec SA policy named AH01, which uses AH protection with SHA1.

```
switch:admin> ipsecconfig --add policy ips sa \
-t AH01 -p ah -auth hmac_sha1
```

3. Create an IPsec SA policy named ESP01, which uses ESP protection with 3DES.

```
switch:admin> ipsecconfig --add policy ips sa \
-t ESP01 -p esp -enc 3des_cbc
```

4. Create an IPsec proposal IPSEC-AHESP to use an AH01 and ESP01 bundle.

```
switch:admin> ipsecconfig --add policy ips sa-proposal \
-t IPSEC-AHESP -sa AH01,ESP01
```

5. Import the preshared key file (e.g., ipseckey.psk) using the **secCertUtil import** command.

6. Create an IKE policy for the remote peer.

```
switch:admin> ipsecconfig --add policy ike -t IKE01 \
-remote 10.33.69.132 -id 10.33.74.13 \
-remoteid 10.33.69.132 -enc 3des_cbc \
-hash hmac_md5 -prf hmac_md5 \
-auth psk -dh modp1024 -psk ipseckey.psk
```

7. Create an IPsec transform TRANSFORM01 configured with transport mode to protect traffic identified for IPsec protection and use IKE01 as a key management policy.

```
switch:admin> ipsecconfig --add policy ips transform \
-t TRANSFORM01 -mode transport -sa-proposal IPSEC \
-AHESP -action protect -ike IKE01
```

8. Create traffic selectors to protect outbound and inbound traffic.

```
switch:admin> ipsecconfig --add policy ips selector \
-t SELECTOR-OUT -d out -l 10.33.74.13 -r 10.33.69.132 \
-transform TRANSFORM01
```

```
switch:admin> ipsecconfig --add policy ips selector \
-t SELECTOR-IN -d in -l 10.33.69.132 -r 10.33.74.13 \
-transform TRANSFORM01
```

9. Verify the IPsec SAs using **ipSecConfig --show manual-sa -a**. Refer to the "IPsec display commands" section for an example.

10. Perform the equivalent steps on the remote peer to complete the IPsec configuration. Refer to your server administration guide for instructions.

Example 4

The following example illustrates how to secure traffic between two systems using protection with MD5 and Manually keyed SAs. The two systems are a switch, the BROCADE300 (IPv4 address 10.33.74.13), and an external UNIX host (IPv4 address 10.33.69.132).

1. On the system console, log in to the switch as Admin and enable IPsec.

```
switch:admin> ipsecconfig --enable
```

2. Create an IPsec Manual SA that uses AH protection with MD5 for outbound traffic:

```
switch:admin> ipsecconfig --add manual-sa -spi 0x300 \
  -l 10.33.74.13 -r 10.33.69.132 -p any -d out -m transport \
  -ipsec ah -ac protect -auth hmac_md5 \
  -auth-key "TAHITEST89ABCDEF"
```

3. Create an SA for inbound traffic.

```
switch:admin> ipsecconfig --add manual-sa -spi 0x200 \
  -l 10.33.69.132 -r 10.33.74.13 -p any -d in \
  -m transport -ipsec ah -ac protect -auth hmac_md5 \
  -auth-key "TAHITEST89ABCDEF"
```

4. Verify the SAs using `ipsecConfig --show manual-sa -a`. Refer to the IPsec display commands section for an example.
5. Perform the equivalent steps on the remote peer to complete the IPsec configuration. Refer to your server administration guide for instructions.

IPsec display commands

To display the IPsec IKE Policy:

```
switch:admin> ipsecconfig --show policy ike -a
IKE-01  version:ikev2 remote:10.33.69.132
        local-id:10.33.74.13  remote-id:10.33.69.132
        encryption algorithm: 3des_cbc
        hash algorithm: hmac_md5
        prf algorithm: hmac_md5
        dh group: 2 1
        auth method:rsasig
        public-key:"/etc/fabos/certs/sw0/thawkcert.pem"
        private-key:"/etc/fabos/certs/sw0/thawkkey.pem"
        peer-public-key:"/etc/fabos/certs/sw0/spiritcert.pem"
```

To display the outbound and inbound SAs in the kernel SA database:

```
switch:admin> ipsecconfig --show manual-sa -a
10.33.69.132[0] 10.33.74.13[0]
  ah mode=transport spi=34560190(0x020f58be) reqid=0(0x00000000)
  A: hmac-md5 7e5aeb47 e0433649 c1373625 34a64ece
  seq=0x00000000 replay=32 flags=0x00000000 state=mature
  created: Oct 15 23:34:55 2008  current: Oct 15 23:35:06 2008
  diff: 11(s)  hard: 2621440(s)  soft: 2100388(s)
  last: Oct 15 23:34:56 2008  hard: 0(s)  soft: 0(s)
  current: 256(bytes)  hard: 0(bytes)  soft: 0(bytes)
  allocated: 4  hard: 0 soft: 0
  sadb_seq=1 pid=10954 refcnt=0
```



```

10.33.74.13[0] 10.33.69.132[0]
  ah mode=transport spi=48095089(0x02ddf71) reqid=0(0x00000000)
  A: hmac-md5 c84d27e5 960d116c bf7c0e4a b232c49e
  seq=0x00000000 replay=32 flags=0x00000000 state=mature
  created: Oct 15 23:34:55 2008   current: Oct 15 23:35:06 2008
  diff: 11(s)   hard: 2621440(s)   soft: 2137448(s)
  last: Oct 15 23:34:55 2008   hard: 0(s)   soft: 0(s)
  current: 540(bytes)   hard: 0(bytes)   soft: 0(bytes)
  allocated: 5   hard: 0 soft: 0
  sadb_seq=0 pid=10954 refcnt=0

```

To display a specified IPsec SA:

```

switch:admin> ipsecconfig --show policy ips \
  sa -t sa-esp-1
sa-esp-1   ipsec-protocol:esp
           encryption algorithm: aes128_cbc 3des_cbc
           authentication algorithm: hmac_sha1 hmac_md5

```

To display all IPsec SA proposals:

```

switch:admin> ipsecconfig --show policy ips \
  sa-proposal -a
ipsec-esp-a-b   SA(s) used:sa-esp-1 sa-ah-1
                lifetime in seconds:infinite
                lifetime in bytes:infinite

ipsec-esp-def   SA(s) used:sa-esp-1
                lifetime in seconds:infinite
                lifetime in bytes:infinite

```

To display all IPsec transforms:

```

switch:admin> ipsecconfig --show policy ips transform -a
policy-A-B   action:auto_ipsec mode:transport
             local:10.33.69.132 remote:10.33.74.13
             sa-proposal:ipsec-esp-a-b
             ike-policy:remote-B

```

To display all IPsec traffic selectors:

```

switch:admin> ipsecconfig --show policy ips selector -a
slt-A-B-any   local:10.33.69.132 remote:10.33.74.13
              direction:outbound upper-layer-protocol:any
              transform-used:policy-A-B

slt-B-A-any   local:10.33.74.13 remote:10.33.69.132
              direction:inbound upper-layer-protocol:any
              transform-used:policy-A-B

```

Using the help command

To use the **--help** command with arguments to display the syntax of specific types and subtypes:

```

switch:admin> ipsecconfig --help add policy ips selector
Usage: ipsecConfig --add policy ips selector ARGUMENTS

ARGUMENTS
  -tag <name>           selector name

```

```

-direction <in|out>      traffic flow direction
-local <addr>           source IPv4 or IPv6 address
-remote <addr>          peers IPv4 or IPv6 address
-transform <name>       transform name
[-protocol <name>]     protocol nam

```

```
switch:admin> ipsecconfig --help modify policy ike
```

```
Usage: ipsecConfig --modify policy ike ARGUMENTS
```

ARGUMENTS

```

-tag <name>             ike policy name
-remote <addr>          peers ipaddress
-id <identifier>        local identifier
-remoteid <identifier> peers identifier
-enc <ALGORITHM>       encryption algorithm
-hash <ALGORITHM>      hash algorithm
-prf <ALGORITHM>       prf algorithm
-dh <number>           dh group number

```

ENCRYPTION ALGORITHM

```
3des_cbc, aes128_cbc, aes256_cbc, null_enc
```

HASH ALGORITHM

```
hmac_md5, hmac_sha1, aes_xcbc
```

PRF ALGORITHM

```
hmac_md5, hmac_sha1, aes_xcbc
```

DH-GROUP

```
modp768(1), modp1024(2), modp2048(14), modp8192(18)
```

References

- [DISR] "DOD IPv6 Standards Profiles for IPv6 Capable Products", v2.0 (6/15/2007).
- [NIST] "A Profile for IPv6 in the U.S. Government", v1.0 (1/31/2007).
- [4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301 (Dec 2005).
- [4302] Kent, S., "IP Authentication Header", RFC4302 (Dec 2005).
- [4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC4303 (Dec 2005).
- [4306] Kaufman, C., Ed., "The Internet Key Exchange (IKEv2) Protocol", RFC 4306 (Dec 2005).

SEE ALSO

None

isiShow

Displays interswitch link (ISL) information.

Synopsis **isishow**

Description Use this command to display the current connections and status of the interswitch link (ISL) for each port on a switch. The command output includes the following information:

- Node world wide name (WWN)
- Domain ID
- Switch name
- ISL connection speed, if applicable
- Bandwidth
- Trunking enabled, if applicable
- QOS - QoS enabled, if applicable
- ENCRYPT - Encryption enabled, if applicable
- COMPRESS - Compression enabled, if applicable
- CR_RECOV - Credit recovery enabled, if applicable
- FEC -Forward Error Correction enabled, if applicable

When issued on a switch that is part of a logical fabric configuration, the **isiShow** command displays logical interswitch links (LISLs) along with regular ISLs. However, speed (sp) displays N/A for logical ports. The bandwidth (bw) displayed is the sum of the bandwidth of all extended ISLs (XISLs) that form the LISL. A shared ISL (XISL) connects the base switches and is shared by different logical fabrics. It allows devices to communicate with each other within the logical fabric.

Connection speed is not applicable to LE_Ports or VE_Ports. For these port types, speed displays as "sp:-----".

Beginning with Fabric OS 7.1.0, this command will display the neighbor WWN information even when the ISL is segmented during exchange link parameter (ELP) or post ELP segmentation phase. In a fabric that has switches running Fabric OS 7.1.0 and pre-7.0.1 firmware, the neighbor information of the segmented links is displayed only on switches running Fabric OS 7.1.0 firmware. In a rare situation when ELP is not exchanged between two switches, the neighbor WWN information will be displayed only on the responder side that received ELP.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To execute **islshow** in a base fabric:

```
switch:user> islshow
1: 2->300 10:00:00:05:1e:43:00:00 100 DCX \
   sp: 8.000G bw: 32.000G TRUNK QOS
2: 8-> 3 10:00:00:05:1e:41:8a:d5 30 B5300\
   sp: 4.000G bw: 16.000G TRUNK QOS
3: 19-> 10 10:00:00:05:1e:41:43:ac 50 B300 sp:\
   8.000G bw: 64.000G TRUNK
```

To execute **islshow** in a logical fabric:

```
switch:user> islshow
1: 33-> 29 10:00:00:60:69:80:4f:84 3 B3200 \
   sp: 2.000G bw: 4.000G TRUNK
2: 39-> 7 10:00:00:60:69:45:68:04 4 B3850 \
   sp: 2.000G bw: 8.000G TRUNK
3: 41-> (incompatible)
4: 47-> (incompatible)
5: 95-> 0 10:00:00:05:1e:01:0b:4a 15 B5100 \
   sp: 2.000G bw: 2.000G TRUNK
6:162->160 10:00:00:60:69:e2:09:fa 5 B2400 \
   sp:----- bw: 0.001G
7:384-> 16 10:00:00:05:1e:37:02:73 39 DCX \
   sp:----- bw: 10.000G
8:385-> 16 10:00:00:05:1e:37:02:73 43 B53000 \
   sp:----- bw: 11.000G
```

To display interswitch links with encryption or compression enabled:

```
switch:user> islshow
1: 33-> 29 10:00:00:60:69:80:4f:84 3 B3200 \
   sp: 2.000G bw: 4.000G TRUNK ENCRYPT
2: 39-> 7 10:00:00:60:69:45:68:04 4 B3850 \
   sp: 2.000G bw: 8.000G ENCRYPT COMPRESS
2: 38-> 6 10:00:00:60:69:45:68:03 4 B3850 \
   sp: 2.000G bw: 8.000G COMPRESS
```

To display interswitch links with Credit Recovery and Forward Error correction enabled:

```
switch:user> islshow
1: 95-> 26 10:00:00:05:33:7e:69:c4 1 sw0 \
   sp: 16.000G bw: 16.000G CR_RECOV FEC
2:164-> 28 10:00:00:05:33:7e:69:c4 1 sw0 \
   sp: 16.000G bw: 16.000G TRUNK QOS CR_RECOV FEC
3:165-> 30 10:00:00:05:33:7e:69:c4 1 sw0 \
   sp: 16.000G bw: 16.000G QOS CR_RECOV FEC
4:166-> 31 10:00:00:05:33:7e:69:c4 1 sw0 \
   sp: 4.000G bw: 4.000G QOS CR_RECOV
5:167-> 29 10:00:00:05:33:7e:69:c4 1 sw0 \
   sp: 16.000G bw: 16.000G CR_RECOV FEC
6:340-> 27 10:00:00:05:33:7e:69:c4 1 sw0 \
   sp: 16.000G bw: 16.000G QOS CR_RECOV FEC
```

```

7:341-> 25 10:00:00:05:33:7e:69:c4 1 sw0 \
      sp: 16.000G bw: 16.000G QOS CR_RECOV FEC
8:342-> 32 10:00:00:05:33:7e:69:c4 1 sw0 \
      sp: 4.000G bw: 4.000G QOS CR_RECOV
9:343-> 24 10:00:00:05:33:7e:69:c4 1 sw0 \
      sp: 16.000G bw: 16.000G CR_RECOV FEC

```

To display interswitch links with D_Port enabled:

```

switch:user> isIshow
  1: 24->343 10:00:00:05:1e:e5:e4:00 1 D-Port_DCX      sp: 16.000G bw: 16.000G
CR_RECOV FEC
  2: 25->341 10:00:00:05:1e:e5:e4:00 1 D-Port_DCX      sp: 16.000G bw: 16.000G
CR_RECOV FEC
  3: 26-> 95 10:00:00:05:1e:e5:e4:00 1 D-Port_DCX      sp: 16.000G bw: 16.000G
CR_RECOV FEC
  4: 27->340 10:00:00:05:1e:e5:e4:00 1 D-Port_DCX      sp: 16.000G bw: 16.000G
CR_RECOV FEC
  5: 29->166 10:00:00:05:1e:e5:e4:00 1 D-Port_DCX      sp: 16.000G bw: 16.000G
CR_RECOV FEC
  6: 30->165 10:00:00:05:1e:e5:e4:00 1 D-Port_DCX      sp: 16.000G bw: 16.000G QOS
CR_RECOV FEC

```

To display the neighbor switch WWN for the segmented ISLs during ELP and post ELP phase:

```

switch:user> isIshow
[...]
1:9-> 2 10:00:00:05:1e:a3:00:59 (incompatible)
[...]

```

SEE ALSO [switchShow](#), [trunkShow](#)

itemList

Lists parameter syntax information.

Synopsis *item_list = element | element white item_list*

element = item | item - item

item = num | slot [white]/ [white] num

slot = num

num = hex | int

int = int digit | digit

hex = 0x hex digit | hex hex digit

digit = 0|1|2|3|4|5|6|7|8|9

hex digit = digit |A|B|C|D|E|F|a|b|c|d|e|f

*white = *["\t\f\r,"]*

Description

All kernel diagnostics have at least one item list parameter to specify which ports to test. The normal default value for this parameter is to select everything.

This is not a command; rather, it is a common parameter to many commands.

If you want to restrict the items to be tested to a smaller set, the parameter value is an item list with the following characteristics:

- It is a comma-separated list of items.
- Each item in the list can be a single element or a range of elements separated by a dash character or a combination of both. For example, "0,3,4-6,1", "0,1,3,4,5,6", and "0 3 4 - 6 1" each select items 0, 1, 3, 4, 5, 6, and 7.
- Spaces and tab stops are skipped.
- Each item might be preceded by an optional slot number followed by a slash ("/).

Besides the syntax rules, there are also some grammatical restrictions on the slot numbers:

- Once specified, a slot selection applies to all items to the right of the slot selections until the next slot selection or the end of the item list. For example, "1/0 - 15" and "1/0 - 1/15" are equivalent.
- If no slot number is specified, user port lists are specified by area number. For instance, "0, 16, 32" and "1/0, 2/0, 3/0" specify the same ports on a 16-port/blade system. On that same system, "1/0, 16, 32" is not a valid list: even though it is legal syntax, the ports do not exist.
- If no slot number is specified, all lists except user port lists use the default slot 0.
- No list type except for user port lists may specify multiple conflicting slot numbers. For instance, "1/0, 2/0, 3/0" is a valid user port list but is not valid for any other type of list.

In the case of conflicting settings within a single item list, an error is generated, as described earlier. In the case of multiple item list parameters, the last one on the command line overrides previous settings.

The exact type of list varies, depending on the test and the parameter; however, the most common are blade ports and user ports. A list of blade ports is most commonly used by ASIC-level tests such as **turboRamTest** and represents which ports on the current blade (specified with `--slot number`) are tested. A list of user ports is used by higher-level tests to specify which user-accessible external ports within the current switch (selected during Telnet login) are tested. When specified in an item list, user ports might be specified by either the area portion of the ports Fibre Channel address or with `slot/port` notation. For nonblade systems, the port number on the silkscreen is the area number, so the two notations are identical.

For item list parameters, the parameter type is PT_LIST and the list type is one of the following:

| Type | Grouping | Description |
|---------|----------|--|
| BPORTS | Blade | Blade ports, internal and external ports. |
| UPOINTS | Switch | User ports, ports with external connections. |
| QUADS | Blade | Quadrants, group of (normally 4) ports. |
| CHIPS | Blade | Chips, Asics within a blade. |
| MINIS | Blade | Mini switches. |
| SLOTS | Chassis | Slots. |
| INDEX | N/A | Anything. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples None

SEE ALSO [portLedTest](#), [portLoopbackTest](#)

killTelnet

Terminates an open Telnet session.

Synopsis `killtelnet`

Description Use this command to terminate an open Telnet session. The command lists all current Telnet and serial port login sessions and information such as session number, login name, idle time, IP address of the connection, and timestamp of when the login session was opened. The command prompts you to specify the number of the session you want to terminate. The list of open sessions displayed with **killTelnet** includes your current session; be sure not kill your own Telnet session.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Examples To terminate an open Telnet connection:

```
switch:admin> killtelnet
Collecting login information...Done
      List of telnet sessions (3 found)
-----
Session No  USER      TTY      IDLE      FROM      LOGIN@
-----
      0      root0     ttyS0    1:17m    -          5:13pm
      1      admin0    pts/0    16.00s   192.168.130.29  6:29pm
      2      admin0    pts/1    3.00s    192.168.130.29  6:31pm
-----
Enter Session Number to terminate (q to quit) 1
Collecting process information... Done.
  You have opted to terminate the telnet session:-
    logged in as "admin0 ", from "192.168.130.29 "
    since " 6:29pm" and has been inactive for "16.00s ",
    the current command executed being: "-rbash ".
    The device entry is: "pts/0 ".
  This action will effectively kill these process(es):-
      USER      PID ACCESS COMMAND
/dev/pts/0      root      12868 f.... login
                root      12869 f.... login
                root      12877 f.... rbash
Please Ensure (Y/[N]): y
killing session... Done!
Collecting login information...Done
      List of telnet sessions (2 found)
-----
Session No  USER      TTY      IDLE      FROM      LOGIN@
-----
      0      root0     ttyS0    1:17m    -          5:13pm
      1      admin0    pts/1    3.00s    192.168.130.29  6:31pm
-----
Enter Session Number to terminate (q to quit) q
```

SEE ALSO `None`

IdapCfg

Maps LDAP AD server roles to default switch roles.

Synopsis `Idapcfg --maprole Idaprole switchrole`

`Idapcfg --unmaprole Idaprole`

`Idapcfg --show`

`Idapcfg --help`

Description Use this command to map a Lightweight Directory Access Protocol (LDAP) Active Directory (AD) server role to one of the default roles available on a switch. This command also provides an option to remove an existing mapping.

This command creates an alias for a customer-defined group which allows a user belonging to that group to login to the switch with the permissions associated with the mapped switch role.

This command supports one-to-one role mapping only. For example, you might map the "SAN administrator" role on the AD server to the "admin" role on the switch, or the "SAN maintenance" role to the switch "operator" role. But the command fails if you attempt to map an already mapped AD server role.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command takes as input an action and its associated arguments. When no operand is specified, the command prints the usage.

This command has the following operands:

--maprole

Maps an LDAP role to a specified switch role. The following operands are required:

Idaprole

Specifies the LDAP role to be mapped to a switch role. The role must be a valid AD server role.

switchrole

Specifies the switch role to which the LDAP role is mapped. Valid switch roles include the following:

- admin
- user
- switchadmin
- zoneadmin
- fabricadmin
- basicswitchadmin
- operator
- securityadmin

--unmaprole

Removes the mapping between an LDAP role and a switch role. Use the **--show** option for a listing of existing mappings. The following operand is required:

ldaprole

Specifies the LDAP AD sever role to be removed from the mapping.

--show

Displays a table of existing mappings between LDAP roles and their corresponding switch role.

--help

Displays the command usage.

Examples

To display current LDAP and switch role map:

```
switch:admin> ldapcfg --show
LDAP Role      |          Switch Role
-----|-----
ldapadmin      |          admin
ldapuser       |          user
SANfabadmin    |          fabricadmin
SANzoneadmin   |          zoneadmin
SANoperator    |          operator
LDAPSANsecadm  |          securityadmin
SANuser        |          user
SAN01secadmin  |          securityadmin
LD_02zoneadmin |          zoneadmin
-----|-----
```

To map an LDAP AD server role to the switch role of "operator":

```
switch:admin> ldapcfg --maprole SANoperator operator
LDAP role SANoperator has been successfully mapped.

switch:admin> ldapcfg --unmaprole SANoperator
LDAP role SANoperator has been successfully unmapped.
```

SEE ALSO [aaaConfig](#), [userConfig](#)

IfCfg

Configures and displays logical fabrics.

Synopsis `lfcfg [--show | --showall] -cfg`
`lfcfg [--show | --showall] -lisl [-v]`
`lfcfg [--show -xisl [slot]port | --showall -xisl`
`lfcfg --lislenable`
`lfcfg --help`

Description Use this command to display logical fabric configuration information, to determine the status of logical interswitch links (LISLs), to enable LISLs between logical switches, and to display information about the XISLs and LISLs associated with each XISL.

A logical switch is a partition created on a physical switch that shares the physical resources of the base fabric while functioning as an independent entity in a "virtual" logical fabric. The logical fabric sits on top of a base physical fabric and ties otherwise disconnected logical switches together to share the same connectivity and physical resources. At the same time, the logical fabric provides protocol and management isolation, and each logical fabric is independently scalable.

The display options provided with this command show the logical fabric configuration for a given logical switch context or for a chassis context. Each logical switch displays only the user ports that are configured to be part of that switch instance. The switch context is defined by the fabric ID. The default context is the base logical switch that you are placed in upon login. The default logical switch context is defined by the fabric ID 128. To change the context, use the `setContext` command.

When issued with the `-cfg` option, this command displays the following information:

Chassis
 Numeric identifier for the chassis.

Chassis WWN
 Chassis world wide name.

Base switch Domain
 The domain ID of the base switch.

For each logical switch, the following information is displayed:

Logical Switch
 Numeric identifier for the logical switch within the chassis.

Base switch
 Yes or No. This field indicates whether or not this logical switch is the base switch.

Fabric Id
 The logical switch fabric ID (FID).

State
 The state of the logical switch: Online or Offline.

| | |
|---|---|
| Switch WWN | The logical switch world wide name. |
| When issued with the -lisl option, the command displays the following information: | |
| FID | Fabric ID of the logical switch. |
| Port | Number of the logical LISL port. |
| remote-domain | Domain ID of the base switch in the remote chassis. |
| Name | Switch name. |
| State | Port state: Online or Offline. |

Associated physical ports

Physical ports associated with the LISL ports.

When **IfCfg** is issued within a logical switch context, only the configuration regarding that switch and the fabrics reachable from that switch is displayed. When the command is issued in a chassis context the information for all chassis in the base fabric reachable from the current chassis is displayed. Executing chassis-level commands requires chassis permissions. Refer to the **userConfig** command for information on setting chassis user permissions.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported only on hardware platforms that are Virtual Fabric-aware and run Fabric OS v6.2.0 or later. Refer to the *Fabric OS Administrator's Guide* for specific hardware support.

You cannot use the **portEnable** command on logical ports. Use **IfCfg** with the **--lislenable** option to re-enable disabled LISL ports on a logical switch.

Operands

This command has the following operands:

| | |
|------------------------|--|
| --show -cfg | Displays information for the fabric ID set by the context in all chassis reachable from the base fabric. |
| --showall -cfg | Displays information for all fabrics in all chassis reachable from the chassis context in which the command is executed. This option requires chassis permissions. |
| --show -lisl | Displays status information about the LISLs in the logical switch set by the context. |
| --showall -lisl | Displays status information of all LISLs in the chassis. This option requires chassis permissions. |

- v** Displays, in addition to the LISLs, the physical ports on the base switch that are associated with the LISL ports. This operand is optional with the **-lisl** option.
- show -xisl [slot/port]** Displays the XISL and the LISLs associated with it for the specified XISL port. For each LISL port, the output displays the FID, LISL State (online/offline) and the local and remote logical switch WWNs. This command must be executed from the base switch.
- showall -lisl** Displays the XISL connections between two base switches for all XISL ports. This command must be executed from the base switch.
- lislenable** Re-enables all LISLs in the fabric that were disabled because of some conflict or error condition in the fabric. This command provides the option of manually reestablishing the LISLs after the error condition has been resolved.
- help** Displays the command usage.

Examples To display logical fabric information for FID 2 in all chassis reachable from the base fabric.

```
switch:admin> ifcfg --show -cfg

----- Chassis: 1 -----
Chassis WWN: 10:00:00:05:1e:39:82:64
Number of Partitions: 2
Base switch domain: 1

    Logical switch: 2   Base switch: YES           Fabric Id: 2
    State: Online(1)   Switch WWN: 10:00:00:05:1e:39:81:67

----- Chassis: 2 -----
Chassis WWN: 10:00:00:05:1e:0b:a4:5e
Number of Partitions: 2
Base switch domain: 2

    Logical switch: 2   Base switch: YES           Fabric Id: 2
    State: Online(1)   Switch WWN: 10:00:00:05:1e:0b:a4:41
```

To display information for all fabrics in all chassis reachable from the base fabric:

```
switch:admin> ifcfg --showall -cfg

----- Chassis: 1 -----
Chassis WWN: 10:00:00:05:1e:39:82:64
Number of Partitions: 2
Base switch domain: 1

    Logical switch: 2   Base switch: YES           Fabric Id: 2
    State: Online(1)   Switch WWN: 10:00:00:05:1e:39:81:67

    Logical Switch: 1   Base switch: NO           Fabric Id: 1
```

```
State: Online(1)      Switch WWN: 10:00:00:05:1e:39:81:66
```

```
----- Chassis: 2 -----
Chassis WWN: 10:00:00:05:1e:0b:a4:5e
Number of Partitions: 2
Base switch domain: 2
```

```
Logical switch: 2    Base switch: YES          Fabric Id: 2
State: Online(1)    Switch WWN: 10:00:00:05:1e:0b:a4:41
```

```
Logical Switch: 1    Base switch: NO          Fabric Id: 1
State: Online(1)    Switch WWN: 10:00:00:05:1e:0b:a4:40
```

To display the LISLs in the logical switch:

```
switch:admin> ifcfg --show -lisl
FID      Port#    remote-domai Name    State
2        384     24      sw0    PT Online
```

Displays status information about the LISLs in the logical switch set by the context:

```
switch:admin> ifcfg --show -lisl -v
ID Port# remote-domain Name State Associated Physical Ports
2  384  24      sw0  PT Online 1/29, 2/41, 3/33, 4/24
```

To display information about all LISLs in the chassis:

```
switch:admin> ifcfg --showall -lisl
FID      Port#    remote-domain Name    State
2        384     24      sw0    PT Online
3        385     24      sw0    PT Online
```

To display all XISLs and the LISLs associated with each XISL:

```
switch:admin> ifcfg --showall -xisl

XISL Port No. : 12/30
LISL Pt. FID LISL State Local LS WWN          Remote LS WWN

450  10  PortOnline 10:00:00:05:1e:48:f8:02 10:00:00:05:1e:58:b2:5a
451  20  PortOnline 10:00:00:05:1e:48:f8:03 10:00:00:05:1e:58:b2:5b
452  30  PortOnline 10:00:00:05:1e:48:f8:04 10:00:00:05:1e:5b:69:d5
453  10  PortOnline 10:00:00:05:1e:48:f8:02 10:00:00:05:1e:5b:69:d4
454  30  PortOnline 10:00:00:05:1e:48:f8:04 10:00:00:05:1e:58:bd:6b
455  10  PortOnline 10:00:00:05:1e:48:f8:02 10:00:00:05:1e:58:bd:6a

XISL Port No. : 12/31
LISL Pt. FID LISL State Local LS WWN          Remote LS WWN
448  10  PortOnline 10:00:00:05:1e:48:f8:02 10:00:00:05:1e:0b:87:dd
449  20  PortOnline 10:00:00:05:1e:48:f8:03 10:00:00:05:1e:0b:87:de
```

To display a specific XISL and the LISLs associated with it

```
switch:admin> ifcfg --show -xisl 12/31
```

```
XISL Port No. : 12/31
LISL Pt. FID LISL State Local LS WWN Remote LS WWN
448 10 PortOnline 10:00:00:05:1e:48:f8:02 10:00:00:05:1e:0b:87:dd
449 20 PortOnline 10:00:00:05:1e:48:f8:03 10:00:00:05:1e:0b:87:de
```

SEE ALSO None

licenseAdd

Adds a license key to a switch.

Synopsis `licenseadd license`

Description Use this command to add a license key to a switch.

Some features of the switch and the fabric to which it is connected are optional, licensed products. Without a valid license installed for such products, their services are not available.

A license key is a string of any length consisting of upper- and lowercase letters and numbers. License keys are case-sensitive. The license must be entered exactly as issued. The system may accept an incorrectly entered license, but the licensed products will not function. After entering the license, use the **licenseShow** command to validate the product associated with the license. If no licensed products are shown, the license is invalid.

After you enter a license, the licensed product is generally available immediately without requiring further action. The following exceptions apply:

- Some licenses may require you to refresh the ports before to activate the license. Depending on your system, use the **portDisable/portEnable**, **switchDisable/switchEnable** or **chassisDisable/chassisEnable** commands to refresh the ports.
- Some licenses may require that you reboot the switch to activate the license. The **licenseAdd** command will prompt you to reboot the switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

license Specifies the license key to be installed. This operand is required.

Examples To add a license key to the switch:

```
switch:admin> licenseadd DXXtN3LmRSMWCSW3XmfSBPfrWKLZ3HMTN73rP9GANJMA
adding license-key [DXXtN3LmRSMWCSW3XmfSBPfrWKLZ3HMTN73rP9GANJMA]
```

SEE ALSO [licenseRemove](#), [licenseShow](#)

licenseIdShow

Displays the system license ID.

Synopsis `licenseidshow`

Description Use this command to display the license ID of the system.

Some features of the switch and the fabric are optional, licensed products. Without a license installed for such products, the services provided by these features are not available.

This command displays the system license ID used for generating and validating licenses on the system. The license ID format consists of eight pairs of hexadecimal values, separated by colons. Each hexadecimal value is between 00 (0) and FF (255).

Notes While the format of this identifier might be similar or even identical to other identifiers in the system, no inferences should be made about the relationships between them as they are subject to change independently of one another.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the license ID:

```
switch:admin> licenseidshow  
a4:f8:69:33:22:00:ea:18
```

SEE ALSO [licenseAdd](#), [licensePort](#), [licenseRemove](#), [licenseShow](#), [licenseSlotCfg](#)

licensePort

Manages Dynamic Ports On Demand (DPOD) licenses.

Synopsis

```
licenseport --release port1[-port2]
licenseport --reserve port1[-port2]
licenseport --show
licenseport --method dynamic | static
```

Description Use this command to manage and display Dynamic Ports on Demand (DPOD) license assignments.

Dynamic Ports On Demand (DPOD) is an optional feature available on all embedded platforms. DPOD takes the expansion capability of static Ports On Demand (POD) and adds the flexibility of activating any available port as long as a valid license is available. In Static mode, POD allows only specific fixed ports to be activated or licensed. With DPOD, any physical port can be made active as long as the total number of licenses is not exceeded.

The Dynamic POD feature assigns ports to the POD license in the order in which they come online until they equal the number of online licensed ports. This command provides the mechanism to make adjustments to the dynamic assignments by reserving assignments for specific ports in the event that there are more online ports than the purchased POD licenses can support.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

On unsupported platforms, this command returns a "not supported" message.

Operands This command has the following operands:

| | |
|--------------------------------|--|
| <i>port1</i> [- <i>port2</i>] | Specifies a single port or a range of ports to which to assign or from which to remove a POD license. A port must be specified when releasing or reserving a port. |
| --release | Releases a license assignment from the specified ports when the switch is using the Dynamic POD method. The ports must be offline for this command to succeed. |
| --reserve | Reserves a license assignment for the specified ports when the switch is using the Dynamic POD method. The ports must be offline for this command to succeed. |
| --show | Displays the POD license assignments. |
| --method | Selects the POD method as one of the following: |
| dynamic | Selects the dynamic POD method. |

static

Selects the static POD method.

The **--method** operand is not supported on Brocade G620 switch.

Examples

To activate Dynamic Ports On Demand:

```
switch:admin> licenseport --method dynamic
The POD method has been changed to dynamic.
Please reboot the switch now for this change to take effect.
```

To release a port from a Dynamic POD license assignment and to display the assignments:

```
switch:admin> portdisable 22
switch:admin> licenseport --release 22
switch:admin> licenseport --show
24 ports are available in this switch
  1 POD license is installed
    Dynamic POD method is in use
24 port assignments are provisioned for use in this switch:
  12 port assignments are provisioned by the base switch license
  12 port assignments are provisioned by the first POD license
* 0 more assignments are added if the second POD license
  is installed
23 ports are assigned to installed licenses:
  12 ports are assigned to the base switch license
  11 ports are assigned to the first POD license
Ports assigned to the base switch license:
  1, 2, 3, 5, 6, 7, 8, 10, 11, 14, 15, 19
Ports assigned to the first POD license:
  0, 4, 9, 12, 13, 16, 17, 18, 20, 21, 23
Ports assigned to the second POD license:
  None
Ports not assigned to a license:
  22
  1 license reservation is still available for use by unassigned ports.
```

To reserve a Dynamic POD license assignment for a range of ports and to display the assignments:

```
switch:admin> portdisable 5-10
switch:admin> licenseport --reserve 5-10
switch:admin> licenseport --show
48 ports are available in this switch
  Full POD license is installed
  Dynamic POD method is in use
48 port assignments are provisioned for use in this switch:
  24 port assignments are provisioned by the base switch license
  24 port assignments are provisioned by a full POD license
  9 ports are assigned to installed licenses:
    9 ports are assigned to the base switch license
    0 ports are assigned to the full POD license
Ports assigned to the base switch license:
  4, 5*, 6*, 7*, 8*, 9*, 10*, 16, 31
Ports assigned to the full POD license:
  None
Ports not assigned to a license:
  0, 1, 2, 3, 11, 12, 13, 14, 15, 17, 18, 19, 20, 21, 22, 23
```

licensePort

24, 25, 26, 27, 28, 29, 30, 32, 33, 34, 35, 36, 37, 38, 39, 40
41, 42, 43, 44, 45, 46, 47

39 license reservations are still available for use by unassigned ports
6 license assignments are held by offline ports (indicated by *)

To disable Dynamic Ports On Demand:

```
switch:admin> licenseport --method static  
The POD method has been changed to static.  
Please reboot the switch now for this change to take effect.
```

SEE ALSO [licenseAdd](#), [licenseRemove](#), [licenseShow](#), [licenseShow](#)

licenseRemove

Removes or deactivates a license key.

Synopsis `licenseremove license`

Description Use this command to remove an existing license key from a switch or to deactivate the license key. The existing license key must be entered exactly as shown by the **licenseShow** command. License keys are case-sensitive.

When the key has been removed, use the **licenseShow** command to verify that the key and the associated product have been uninstalled. You must reboot the switch after removing a license. For a switch that has no licenses installed, **licenseShow** displays "No licenses."

Upgradeable licenses, such as slot-based licenses or Universal Time-based licenses cannot be permanently removed; they remain in the database but are not displayed in the **licenseShow** output.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operand is required:

license Specifies the license key to be removed or deactivated. This operand is required.

Examples To remove a license key from the switch:

```
switch:admin> licenseremove bQebzbRdScRfc0iK
removing license key [bQebzbRdScRfc0iK]
```

To deactivate a slot-based license:

```
switch:admin> licenseremove \
DXXtN3LmRSMWCSW3XmfSBPfrWKLZ3HMTN73rP9GANJMA \
removing license-key \
[DXXtN3LmRSMWCSW3XmfSBPfrWKLZ3HMTN73rP9GANJMA]
```

```
switch:admin> licenseshow
SSezcSec9RXTf0dj:
Performance Monitor license
KgrfCBHgrtFBrGKH4D7f9S3FBX7K3MtTtBHKrYHA4CMB:
A7N9rATZLYgFa7JBfmGEJKETgCMHFGQMY4gfLmGAA4GA:
Advanced FICON Acceleration (FTR_AFA) license
Capacity 6
Consumed 2
Configured Blade Slots 3,4
KSYFYAtAfPGHDRJfCYQrC4Q9T9CYYmXDBJgLB:
Enhanced Group Management license
XFXYHmPtCWC93CLgBD9BZD9AmTFgDStFDJGmaLKC9FgWafgSE:
8 Gig FC license
```

licenseRemove

```
Expiry Date 03/06/2009
License is expired
Hf7MBEEGCFNmTAWXXF99RtHXQN4RRtM3mLGtrWZLAMaTaAY9EB:
Storage Application Services license
Expiry Date 03/07/2009
```

SEE ALSO [licenseAdd](#), [licenseIdShow](#), [licenseShow](#)

licenseShow

Displays current license keys.

Synopsis `licenseshow`

Description Use this command to display current license keys, along with a list of licensed products enabled by these keys. Depending on the type of license, this command displays the following information:

Permanent licenses

- License key
- Associated product

Temporary and universal time-based licenses

- License key
- Associated product
- Expiration date or expiration notice if the license has expired

Slot-based licenses

- License key
- Associated product
- Capacity (number of slots purchased)
- Consumed (number of slots configured to use the license)
- Configured Blade Slot Positions (slot numbers of the configured blade slots)

When no licenses are installed, the message "No license installed on this switch" is displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

On the Brocade DCX and DCX 8510 chassis family running Fabric OS v7.0.0 and later, the interchassis link (ICL) license is displayed as an ICL Ports on Demand (POD) license. The first ICL POD license indicates half of the ICL ports; the second ICL POD license indicates all of the ICL ports. When you upgrade a DCX to Fabric OS v7.0.0, the license display does not change. Refer to the example section for an illustration.

Operands None

Examples To display the license keys on a switch with permanent licenses installed:

```
switch:admin> licenseshow
S9bddb9SQbTAceeC:
    Fabric license
eezeRRySff0fSe:
    Remote Switch license
bzbzRcbcSc0c0SY:
    Remote Fabric license
dSeR9RcSeeTfSAq:
    Extended Fabric license
```

licenseShow

```
RyeSzRScycTzft09:
    Entry Fabric license
RyeSzRScyczft0G:
    Trunking license
RyeSzRScycS0ft09:
    4 Domain Fabric license
```

To display the license keys on a switch with temporary (expired) licenses installed:

```
switch:admin> licenseShow
7QmYFYJrmDgE9tTS4AYXB9trYSGtMtrQZSTK4ZSC7FC9ZAYAgE:
    Integrated Routing license
    Expiry Date 01/16/2008
    License is expired
33YBfZfKZ3tQKrRJRtgmS3JDtCL99P4fYrJYQP7GffS4ASmNE:
    Enterprise Bundle license
    Expiry Date 01/16/2008
    License is expired
```

To display the license keys on a switch with universal time-based licenses:

```
switch:admin> licenseshow
DAmHTPgQ7KDtKrEYQC7X7STF9HJDL7TmTWRmZPmSTSE49AEfaE:
    Trunking license
    Expiry Date 11/11/2008
    License is expired
H47CFsa93aKgKJM9NWMYEMaLrATQWDHCgHfZftWGGgNCYAJaBA:
    High-Performance Extension over FCIP/FC license
    Expiry Date 12/20/2008
```

To display a slot-based 10G FC and 10G Gigabit Ethernet license on a Brocade DCX 8510-8:

```
switch:admin> licenseshow
tKLFTNAPDQtEPHFcf9YSPrafrTJCQALP49fXEGHA4SPB:
    10 Gigabit FCIP/Fibre Channel (FTR_10G) license
    Capacity 3
    Consumed 3
    Configured Blade Slots 1,3,12
```

```
switch:admin> slotshow -m
```

| Slot | Blade Type | ID | Model Name | Status |
|------|------------|----|------------|---------|
| 1 | SW BLADE | 97 | FC16-32 | ENABLED |
| 2 | UNKNOWN | | | VACANT |
| 3 | AP BLADE | 75 | FX8-24 | ENABLED |
| 4 | UNKNOWN | | | VACANT |
| 5 | CORE BLADE | 98 | CR16-8 | ENABLED |
| 6 | CP BLADE | 50 | CP8 | ENABLED |
| 7 | CP BLADE | 50 | CP8 | ENABLED |
| 8 | CORE BLADE | 98 | CR16-8 | ENABLED |
| 9 | UNKNOWN | | | VACANT |
| 10 | UNKNOWN | | | VACANT |
| 11 | UNKNOWN | | | VACANT |
| 12 | SW BLADE | 96 | FC16-48 | ENABLED |

To display an ICL license on a Brocade DCX 8510-8:

```
switch:admin> licenseshow
X3ffNTZM9CNmM4SKFLYTGS4WmCRCgAZZBJDTB:
  Inter Chassis Link (2nd POD) license
```

To display an ICL license on a Brocade DCX before and after a firmware upgrade to Fabric OS v7.0.0.

```
switch:admin> licenseshow
X3ffNTZM9CNmM4SKFLYTGS4WmCRCgAZZBJDTB:
  Inter Chassis Link (16 link) license
```

To display the WAN Rate upgrade licenses on Brocade 7840 switch:

```
switch:admin> licenseshow
gFEfTam4YHDLGXZBZ7RHHMBGmrrS3aNWgANKggEAGAJB:
  WAN Rate Upgrade 1 license
ARSNJE4GC9aJ4Y7FrBgYaEWBfKK3EMRFHSBhtHAEERB:
  WAN Rate Upgrade 2 license
```

SEE ALSO [licenseAdd](#), [licenseIdShow](#), [licenseRemove](#)

licenseSlotCfg

Configures and displays slot-based licensed features.

Synopsis `licenseslotcfg --add feature slot`

`licenseslotcfg --remove feature slot`

`licenseslotcfg --show`

`licenseslotcfg --help`

Description

Use this command to configure and manage licenses for the Brocade FX8-24 extension blade on the slot where the blade is installed.

Slot-based licenses allow you to select the slots the license will enable up to the purchased capacity and thereby increase existing capacity without disrupting the slots for which licensed features are already enabled.

There is a separate slot-based license key for each licensed feature supported on the blade.

The Brocade FX8-24 supports slot-based license keys for each of the following features:

- 10 GbE license - Enables the two 10GbE ports on the Brocade FX8-24.
- Advanced Extension license - Enables FCIP Trunking and Adaptive Rate Limiting.
- Advanced FICON Acceleration license - Accelerates FICON tape read and write and IBM Global Mirror data replication operations over distance.

The Brocade 7840 supports slot-based license keys for each of the following features:

- WAN Rate Upgrade 1 - Provides additional throughput up to 10 Gbps.
- WAN Rate Upgrade 2 - Provides unlimited throughput as supported by the hardware.

A license key with the specified capacity must be installed with the **licenseAdd** command before you can enable the feature on a specified slot with the **licenseSlotCfg** command. Refer to the *Fabric OS Administrator's Guide* for more information.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

feature

Specifies the licensed feature to be added at the specified slot. The feature is specified as a tag from the following set of licenses:

FTR_AE

Advanced Extension license

FTR_AFA

Advanced FICON Acceleration license

FTR_10G

10 Gigabit Ethernet license

| | |
|-----------------|--|
| FTR_UPG1 | WAN Rate Upgrade 1 license |
| FTR_UPG2 | WAN Rate Upgrade 2 license |
| slot | Specifies the slot number for the Brocade FX8-24 extension blade. This number corresponds to the physical blade slot number on the chassis. |
| --add | Adds a slot-based license to the specified slot. |
| --remove | Removes a slot-based license from the specified slot. This operation frees up the license to be assigned to another slot. You must disable the applications that use the license on this slot before you can deactivate the license. |
| --show | Displays slot assignments for all slot-based licenses in the chassis. |
| --help | Displays the command usage. |

Examples

To configure the blade slots 3, 4, 11, and 12 to enable the license on these slots:

```
switch:admin> licenseslotcfg --add FTR_AE 3
Blade slot-3 added to FTR_AE slot-based license configuration
Remaining capacity for FTR_AE slot-based license = 7

switch:admin licenseslotcfg --add FTR_AE 4
    Blade slot-4 added to FTR_AE slot-based license configuration
Remaining capacity for FTR_AE slot-based license = 6

switch:admin> licenseslotcfg --add FTR_AE 11
    Blade slot-11 added to FTR_AE slot-based license configuration
Remaining capacity for FTR_AE slot-based license = 5

switch:admin> licenseslotcfg --add FTR_AE 12
Blade slot-12 added to FTR_AE slot-based license configuration
Remaining capacity for FTR_AE slot-based license = 4
```

To display the enabled licenses:

```
switch:admin> licenseslotcfg --show
FTR_10G license - blade slots configured = 1,3,5,12
FTR_AE license - blade slots configured = 3,4,11,12
FTR_AFA license - blade slots configured = 3,4
```

To deactivate the Advanced Extension license on slots 3 and 12, and to display the results:

```
switch:admin> licenseslotcfg --remove FTR_AE 3
Blade slot-3 removed from FTR_AE slot-based
license configuration

switch:admin> licenseslotcfg --show
FTR_10G license - blade slots configured = 1,3,5,12
FTR_AE license - blade slots configured = 12
```

licenseSlotCfg

```
FTR_AFA license - blade slots configured = 3,4
```

```
switch:admin> licenseslotcfg --remove FTR_AE 12  
Blade slot-12 removed from FTR_AE slot-based license configuration
```

```
switch:admin> licenseslotcfg --show  
FTR_10G license - blade slots configured = 1,3,5,12  
FTR_AFA license - blade slots configured = 3,4
```

To view the WAN Rate Upgrade licenses:

```
switch:admin> licenseslotcfg --show  
licenseslotcfg: Command is not supported on this platform
```

SEE ALSO [licenseAdd](#), [licenseIdShow](#), [licenseRemove](#)

linkCost

Sets or displays the Fabric Shortest Path First (FSPF) cost of a link.

Synopsis `linkcost [[slot/]port [cost]]`

Description Use this command to set or display the cost of an interswitch link (ISL). The cost of a link is a dimensionless positive number. The Fabric Shortest Path First (FSPF) protocol compares the cost of various paths between a source switch and a destination switch by adding the costs of all the ISLs along each path. FSPF chooses the path with minimum cost. If multiple paths exist with the same minimum cost, FSPF distributes the load among these paths. The default link cost value is 500.

When executed without operands, the command displays the current cost of each port on the switch, including non-ISLs. An E_PORT suffix is appended to the interface number of active ISLs. If a static cost is assigned to a port, a STATIC suffix is appended to the link cost. In this case, only the current link cost displays. Use **interfaceShow** to display both the default and current link costs.

Notes This command sets a non-default, "static" cost for any port except EX/VEX ports. Use **fcrRouterPortCost** to configure EX/VEX ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command cannot be executed on a logical ISL (LISL).

Operands This command has the following operands

slot

For bladed systems only, specifies the slot number for which to set or display the cost, followed by a slash (/).

port

Specifies the port number for which to set or display the cost, relative to its slot for bladed systems. Use **switchShow** to list of valid ports.

cost

Specifies the static cost of the link connected to the specified port. Recommended cost values are 0 through 32765. Setting static link cost values above 32765 will cause access problems. When you set a link cost value between 32766 and 65534, some parts of the fabric may become inaccessible. When you set the maximum link cost value of 65535, the port will become unusable. In either case, the command displays an appropriate warning message, and you are prompted to continue or to cancel and reissue the command with a lower link cost value. A value of 0 removes the static cost and the port reverts to its default link cost. If *cost* is not specified, the command displays the current cost of the specified port.

Examples To display the link costs for all ports on a switch:

```
switch:admin> linkcost
      Port                Cost
-----
1/0   (E_PORT)           500
1/1                               500
1/2                               500
1/3                               500
1/4                               500
1/5                               500
1/6                               500
1/7                               500
1/8                               500
1/9   (E_PORT)           500
1/10                              500
1/11                              500
1/12                              500
1/13                              500
1/14                              500
1/15                              500
Type <CR> to continue, Q<CR> to stop:
```

To set the ISL cost on a port:

```
switch:admin> linkcost 1/9 1000
```

To display the new cost value on the same port:

```
switch:admin> linkcost 1/9

Interface1/9 (E_PORT) Cost    1000 (STATIC)
```

To delete the cost value and reset to default:

```
switch:admin> linkcost 1/9 0
```

To display the change:

```
switch:admin> linkcost 1/9

Interface1/9 (E_PORT) Cost    500
```

To set the ISL cost to a value outside of the recommended range:

```
switch:admin> linkcost 1/9 32766
The link cost entered may prevent some parts of the fabric
from being accessible. If you do not want this to happen, choose n|no
and run 'linkcost' again with a value lower than 32766.
Do you want to continue? (yes, y, no, n): [no]y

switch:admin> linkcost 1/9

Interface1/9 (E_PORT) Cost    32766 (STATIC)
```

To set the ISL cost to the maximum value:

```
switch:admin> linkcost 1/9 65535
The link cost entered will cause the port to become unusable.
If you do not want this to happen, choose n|no
and run 'linkcost' again with a value lower than 65535.
Do you want to continue? (yes, y, no, n): [no] n
```

SEE ALSO [interfaceShow](#), [lsDbShow](#), [topologyShow](#), [uRouteShow](#), [fcrRouterPortCost](#)

logicalGroup

Creates and manages groups of monitored elements.

Synopsis

```
logicalgroup --create group_name -type group_type
  [-feature feature_name -pattern value | -members member_list]

logicalgroup --delete group_name [-force]

logicalgroup --addmember group_name -members member_list

logicalgroup --delmember group_name -members member_list

logicalgroup --clone existing_group_name
  -name new_group_name

logicalgroup --update group_name -feature feature_name -pattern value

logicalgroup --restore group_name

logicalgroup --show [group_name] [-details]

logicalgroup --help
```

Description Use this command to create and modify groups of elements that are to be monitored using the same set of thresholds. For example, you can create a group of ports that behave in a similar manner, such as UNIX ports or long-distance ports.

The elements in a group must be the same type: ports, circuits, or SFP transceivers. By creating a group of similar elements, you can manage these elements as a single entity.

The maximum number of user-defined groups is 64 per logical switch.

Notes This command requires a Fabric Vision license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

group_name Specifies the name of the group to be configured. The name for the group must be unique; it is not case-sensitive and can contain up to 32 characters.

member_list

Specifies the elements in the group. The member list must be enclosed in double quotation marks and can consist of the following:

Port and SFP groups:

- A single port, for example, "8" or "2/8" on bladed systems.
- A port range where the beginning and end port are separated by a dash, for example, "8-13" or "2/8-13" on bladed systems.
- A set of ports and port ranges separated by commas, for example, "3,5,7-10,15" or "2/3,2/5,3/7-10,4/15" on bladed systems.

Circuit groups:

- A single circuit in VE_Port/circuit format, for example, "33/5".
- A circuit range in VE_Port/circuit format, for example, "33/1-3".
- A set of circuits separated by commas, for example, "33/1-3,33/5".

--create *group_name*

Creates a group of monitored elements.

-type *group_type*

Specifies the type of elements present in the group. Valid types include the following:

port

The elements are FC ports.

circuit

The elements are FCIP circuits.

sfp

The elements are SFP transceivers.

-feature *feature_name*

Specifies the existing feature name for the group. Valid feature names are PORTNAME and NODEWWN.

-pattern *value*

Specifies the wildcard characters while defining the feature characteristics. The wildcard characters "" for any string, "?" for any single character, "[expr]" for one character from the set specified in the expression, or "!" for negation of the string, are supported. If "!" is specified in the pattern, the pattern must be in single quotes.

-members *member_list*

Specifies the elements that are to be in the group.

--delete *group_name* [-force]

Deletes a logical group of monitoring elements. You cannot delete a predefined group. You cannot delete a group that is used by any rules. The **-force** option overrides the default behavior. If a logical group is present in user-defined rules, the **-force** option deletes all the rules that are configured with the given group and then deletes the group.

--addmember *group_name*

Adds members to the group.

-members *member_list*

Specifies the elements that are to be added to the group.

- delmember *group_name***
Deletes members from the group.
- members *member_list***
Specifies the elements that are to be deleted from the group.
- clone**
Creates a replica of an existing group. The new group has all of the members of the existing group. You can further modify the newly created group. The following operands are required:
- existing_group_name***
Specifies the name of an existing group. The group can be a predefined group or a user-defined group.
- name *new_group_name***
Specifies the name of the group to be created.
- update**
Changes the characteristic string of an existing group.
- feature *feature_name***
Specifies the existing feature name for the group. Valid feature names are PORTNAME and NODEWWN.
- pattern *value***
Specifies the wildcard characters while defining the feature characteristics. The wildcard characters "*" for any string, "?" for any single character, "[expr]" for one character from the set specified in the expression, or "!" for negation of the string, are supported. If "!" is specified in the pattern, the pattern must be in single quotes.
- restore**
Restores the membership entries for the group.
- group_name***
Restores entries for a single specified group. The name can be a predefined or a user-defined group created with feature option.
- show**
Displays detailed information for a single group or all groups in MAPS. For each group, the group name, type, member count, and group members are displayed.
- group_name***
Displays information for a single specified group. The name can be a predefined or a user-defined group.
- [-details]**
Displays detailed information for each group. This operand is optional.
- help**
Displays the command usage.

Examples To create a group with three port members:

```
switch:admin> logicalgroup --create CRITICAL_PORTS
                    -type port -members "2, 10, 22"
```

To add members to a group:

```
switch:admin> logicalgroup --addmember CRITICAL_PORTS
               -members "4, 5"
```

To delete members from a group:

```
switch:admin> logicalgroup --delmember CRITICAL_PORTS
               -delmember "5"
```

To clone a group:

```
switch:admin> logicalgroup --clone CRITICAL_PORTS
               -name LONG_DISTANCE_PORTS
```

To display the members of a group:

```
switch:admin> logicalgroup --show CRITICAL_PORTS
Group Name      |Predefined |Type |Member Count |Members
-----
CRITICAL_PORTS  No         Port  9           2, 4, 10, 22
```

To delete a group:

```
switch:admin> logicalgroup --delete CRITICAL_PORTS
```

To update a group:

```
switch:admin> logicalgroup --update dynGroup -feature portname -pattern "port1"
```

To restore a group:

```
switch:admin> logicalgroup --restore ALL_HOST_PORTS
```

To display detailed information about the group:

```
switch:admin> logicalgroup --show group1 -details
GroupName       : group1
Predefined      : No
Type            : Port
MemberCount     : 6
Members         : 2, 11, 20-23
Added Members   : 11
Deleted Members : 1
Feature         : PORTNAME
Pattern         : port2*
```

SEE ALSO [mapsConfig](#), [mapsDb](#), [mapsPolicy](#), [mapsRule](#), [mapsSam](#)

login

login

Logs in as new user.

Synopsis `login`

Description Use this command to log in to the switch with another user name and password, without first logging out from the original session. If you originally connected through a Telnet or rlogin session, that session is left open.

This command allows you to access commands that you cannot access at your current user level.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To log in as admin from the login user:

```
switch:user> login
login: admin
Password: *****
```

SEE ALSO [logout](#)

logout

Logs out from a shell session.

Synopsis **logout**

Description Use this command to log out from a shell session. Remote login connections are closed and the local serial connections return to the **login** prompt.

The **exit** command is accepted as a synonym for **logout**, as is **Ctrl-D** at the beginning of a line.

Operands None

Examples To log out from an rlogin session:

```
switch:admin> logout
```

SEE ALSO [login](#)

IsanZoneShow

Displays logical SAN zone information.

Synopsis `Isanzoneshow [-s] [-f fabricid] [-w wwn] [-z zonename] [-d | -deviceinfo] [-o | -sort]`

Description Use this command to display the inter-fabric zones or LSAN zones. These zones are normal WWN zones created in FC Router EX_Port-connected fabrics and backbone fabrics. The LSAN zones are identified by the text string "Isan_" in the zone name. Note that the string is case insensitive so "LSAN_" also is valid. The FC Router uses these zones to establish the inter-fabric device import and export policy. The LSAN zones are established by zoning administration in each EX_Port-attached fabric and backbone fabric. Inter-fabric device sharing is allowed between two devices if the LSAN zones defined in their respective fabrics both allow the two devices to communicate; for example, the intersection of LSAN zones in two fabrics define the device sharing policy.

The LSAN zones are listed by fabric. Zone membership information (information about the devices in the zone) is provided for each LSAN zone. The default output displays only WWNs of the zone members.

Search parameters **-f**, **-w**, and **-z** allow searching for LSAN zones based on fabric ID, WWN of an LSAN zone member, or LSAN zone name.

"No LSAN zone found" is displayed if there is no LSAN zone information available at this FC Router.

Each LSAN zone entry displays the following:

Fabric ID

The ID of the fabric in which the LSAN zone was created.

Zone Name

The zone name.

Zone Members

The zone members or devices. The default output displays the WWN of the zone members.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-s *state*

Displays state information for the device and the LSAN zone information. Valid states include the following:

Configured

Device is configured to be in an LSAN, but the device is neither imported nor exists in this fabric.

Initializing

Device is in an intermediate state. It is not yet imported into the fabric.

EXIST

Device exists in this fabric (the fabric of the zone entry).

| | |
|-------------------------|--|
| Imported | Device has been imported (proxy created) into this fabric. |
| -f fabricid | Displays LSAN zones in the specified fabric. |
| -w wwn | Displays LSAN zones containing the specified port WWN. The WWN format is xx:xx:xx:xx:xx:xx:xx:xx . |
| -z zonename | Displays LSAN zones with the specified zone name. The database for zones is displayed per switch, which can differ from the database stored on the other FCR switches. |
| -d -deviceinfo | Displays the fabric IDs of the LSAN devices. This command displays fabric IDs for Imported and EXIST devices. The devices that do not exist in the corresponding edge fabrics are displayed as configured . |
| -o -sort | Displays the WWNs of the zone members in ascending order. |

Examples To display the LSAN zones:

```
switch:admin> lsanzoneshow
Fabric ID: 4 Zone Name: lsan_fcr10_0
50:05:07:65:05:84:0b:83
50:05:07:65:05:84:09:0e
10:00:00:00:c9:2b:6a:68
21:00:00:20:37:18:22:55
Fabric ID: 5 Zone Name: lsan_fcr11_0
10:00:00:00:c9:2b:6a:68
21:00:00:20:37:18:22:55
50:05:07:65:05:84:0b:83
50:05:07:65:05:84:09:0e
switch#
```

To display the LSAN zones with device information:

```
switch:admin> lsanzoneshow -d
Fabric ID: 10 Zone Name: LSAN_10
30:02:00:05:1e:61:23:8f EXIST in FID 10
30:02:01:05:1e:61:23:8f EXIST in FID 10
30:00:00:05:1e:61:23:8f Configured
30:06:00:05:1e:61:23:8f Imported from FID 20
30:06:01:05:1e:61:23:8f Imported from FID 20
Fabric ID: 20 Zone Name: LSAN_20
30:02:00:05:1e:61:23:8f Imported from FID 10
30:02:01:05:1e:61:23:8f Imported from FID 10
30:01:00:05:1e:61:23:8f Configured
30:06:00:05:1e:61:23:8f EXIST in FID 20
30:06:01:05:1e:61:23:8f EXIST in FID 20
```

lsanZoneShow

To display state information for the device and sort the WWNs in ascending order:

```
switch:admin> lsanzoneshow -s -sort
Fabric ID: 12 Zone Name: lsan_zone1
    30:06:00:05:1e:61:23:8f EXIST
    30:0c:00:05:1e:61:23:8f Imported
Fabric ID: 16 Zone Name: lsan_zone1
    30:06:00:05:1e:61:23:8f Imported
    30:0c:00:05:1e:61:23:8f EXIST
Fabric ID: 20 Zone Name: lsan_bb
    10:06:00:01:1e:61:23:8f Configured
    10:06:00:05:1e:61:23:8f Configured
    20:03:00:05:1e:61:23:8f Configured
    20:06:00:05:1e:61:23:8f Configured
```

SEE ALSO [fcrFabricShow](#), [fcrPhyDevShow](#), [fcrProxyDevShow](#), [fcrRouteShow](#), [switchShow](#)

IsCfg

Configures and manages a logical switch

Synopsis

```

IsCfg --create FID [-b | -base] [-f | -force]

IsCfg --delete FID [-f | -force]

IsCfg --config FID [-slot slot1[-slot2]]
  [-port [ port1[-port2]] [-q | -qsfp]
  [-f | -force]

IsCfg --restore_to_default FID

IsCfg --restore_slot_to_default slot

IsCfg --change FID [[-newfid FID] | [-base]] [-force]

IsCfg --show [-ge] [-provision] [-n | -name]

IsCfg --help

```

Description Use this command to create a logical switch and to modify logical switch configurations.

The logical switch feature provides the ability to partition a single physical switch into multiple switch instances. Each of these switch partitions is referred to as a logical switch (LS). The logical switch feature allows you to configure multiple logical fabrics on top of a base (physical) fabric. Each logical fabric is made up of logical switches that share the physical resources of the base fabric, for example, interswitch link (ISL) connectivity. At the same time, protocol and management isolation of each logical fabric is maintained, and each logical fabric can scale independently.

The Default Logical Switch is created by the system and cannot be deleted. All switch ports not explicitly assigned to a logical switch are part of the default logical switch.

The Virtual Fabric (VF) feature must be enabled on the switch before you can configure a logical switch. Use the **fosconfig --enable vf** command to enable the feature. Use the **fosconfig --show** command to determine whether the VF feature is enabled or disabled on the switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Not all commands that support the **-force** option will prompt for user input when used without the **-force** option.

Operands This command has the following operands:

--create

Creates a logical switch instance. The following operands are supported.

FID

Specifies the Fabric ID. Each logical switch in a chassis is assigned a unique fabric identifier. The FID address space is shared between logical switches and EX_Ports. Valid FID values are integers between 1 and 128. The default logical switch is assigned FID 128 by default. This operand is required.

- b | -base**
Creates a base logical switch on the chassis. A base logical switch communicates among different logical switches. Legacy switches can be connected to the base logical switch through EX_Ports, and interswitch links (ISLs) between base logical switches enable communication among different logical switches. This operand is optional.
- f | -force**
Executes the command without confirmation. This operand is optional.
- delete**
Deletes a logical switch with the specified fabric ID. The specified logical switch must exist and no ports should be configured on this partition. You must remove all ports from the logical switch before deleting the logical switch instance. Use the **IsCfg --config** command to remove the ports.
- FID**
Specifies the Fabric ID of the logical switch. This operand is required.
- f | -force**
Executes the command without confirmation. This operand is optional.
- config**
Configures the specified logical switch. This command assigns ports to the logical switch specified by a given FID. The ports are removed from the partition on which they are currently configured. This command prompts for confirmation, indicating that the specified ports will be disabled. The following operands are supported:
- FID**
Specifies the fabric ID of the logical switch. This operand is required.
- slot slot1[-slot2]**
Specifies the slot number or a range of slot numbers separated by a dash, for example **-slot 3-5**. This operand is required.
- port port1[-port2]**
Specifies the ports to be assigned to the logical switch. Provide a valid port, or a range of ports separated by a dash, for example **-port 3-8**. This operand is optional; if omitted, all ports on the specified slots are assigned.
- q | -qsfp**
Moves all ports in a quad small form-factor pluggable (QSFP) group to be assigned to the logical switch. This operand is optional.
- f | -force**
Executes the command without confirmation. This operand is optional.
- restore_to_default FID**
Moves all vacant ports in the logical switch specified by **FID** to the default switch. Use this command when **IsCfg --show** displays no ports, but the switch continues to generate errors indicating that there are ports on the switch.
- restore_slot_to_default slot**
Moves all ports on a specified slot to the default switch.

- change** Changes the fabric ID of a logical switch, creates a base logical switch out of an existing logical switch, or removes base switch properties. The **-newfid** and **-base** operands are exclusive and may not be combined. The following operands are supported:
- FID** Specifies the Fabric ID of the logical switch. This operand is required.
- n | -newfid FID** Changes the fabric ID of an existing logical switch. This command effectively removes the logical switch from a given logical fabric and makes it part of another logical fabric.
- b | -base** Turns an existing logical switch into a base switch. When this command is issued on a switch that is already a base switch, this command removes the base switch properties. This command disables the current logical switch. After making the change, you must re-enable the switch.
- f | -force** Executes the command without confirmation. This operand is optional.
- show** Displays the partition configuration on a switch or chassis. Without any operands, the command displays all logical switches and the FC ports assigned to them. For each switch, the FID, switch role, and domain ID are displayed: base switch (bs) or default switch (ds). The following operands are optional with the **--show** option.
- ge** Displays partition configuration information for GbE ports. This operand is valid only on the Brocade FX8-24 platform.
- provision** Displays the partition configuration for all slots, regardless of the slot's status. This operand is valid only on a chassis and can be used with or without the **-ge** option.
- n | -name** Displays fabric ID, switch type, domain ID, switch name, and fabric name of the logical switch. This operand can be used with the **-n** or **-name** option.
- help** Displays the command usage.

Examples To create a base switch:

```
switch:admin> lscfg --create 1 -base
Creation of a base switch requires that the proposed
  new base switch on this system be disabled.
Would you like to continue [y/n]?: y
About to create switch with fid=1. Please wait...

Switch successfully created.

Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
```

To create a logical switch identified by fabric ID 2:

```
switch:admin> lscfg --create 2
A Logical switch with FID 2 will be created with default configuration.
Would you like to continue [y/n]?: y
About to create switch with fid=2. Please wait...
switch1 Updating flash ...
2013/05/02-10:13:59, [ZONE-1034], 373, SLOT 4 | FID 2, INFO, switch_2, \
  A new zone database file is created.
2013/05/02-10:14:15, [FSSM-1002], 374, SLOT 5 | CHASSIS, INFO, switch, \
  HA State is in sync.
All service instances in sync
2013/05/02-10:14:15, [FSSM-1002], 375, SLOT 4 | CHASSIS, INFO, switch, \
  HA State is in sync.
Logical Switch with FID (2) has been successfully created.

Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
2013/05/02-10:14:17, [PMGR-1001], 376, SLOT 4 | CHASSIS, INFO, switch, \
  Attempt to create switch 2 succeeded.
```

To create a base switch with FID 2 without confirmation:

```
switch:admin> lscfg --create 2 -base -force
About to create switch with fid=2. Please wait...
Switch successfully created.

Logical Switch has been created with default configurations.
Please configure the Logical Switch with appropriate switch
and protocol settings before activating the Logical Switch.
```

To delete a logical switch:

```
switch:admin> lscfg --delete 2
The Logical switch with FID 2 will be deleted.
Would you like to continue [y/n]?: y
2013/05/02-10:18:08, [FW-1424], 387, SLOT 4 | FID 2, WARNING, switch_2, \
  Switch status changed from HEALTHY to MARGINAL.
2013/05/02-10:18:08, [FW-1439], 388, SLOT 4 | FID 2, WARNING, switch_2, \
  Switch status change contributing factor Switch offline.
All active login sessions for FID 2 have been terminated.
Switch successfully deleted.
2013/05/02-10:18:29, [PMGR-1003], 389, SLOT 4 | CHASSIS, INFO, pluto_19, \
  Attempt to delete switch 2 succeeded.
```

To assign ports to a logical switch:

```
switch:admin> lscfg --config 2 -port 10-12
This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Making this configuration change. Please wait...
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

To assign ports to a logical switch without confirmation:

```
switch:admin> lscfg --config 2 -port 0-4 -force
Configuration change successful.
Making this configuration change. Please wait...
Please enable your ports/switch when you are ready to continue.
```

To move all ports in a QSFP group to a logical switch:

```
switch:admin> lscfg --config 10 -slot 5 \
-port 0-20 -qsfp
This operation requires that the affected ports be disabled and will move all \
ports(0-23) in qsfp to fid 10.
Would you like to continue [y/n]?: y

Making this configuration change. Please wait...
Dispatch a request to kernel-land component: swc
Dispatch a request to kernel-land component: swc
Configuration change successful.
Please enable your ports/switch when you are ready to continue.
```

To display the logical switch configuration for :FC ports only:

```
switch:admin> lscfg --show

Created switches FIDs(Domain IDs): 128(ds)(10) 1(bs)(1) 2(1)
Port 0 1 2 3 4 5 6 7 8 9
-----
FID 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 2 | 2
-----
Port 10 11 12 13 14 15 16 17 18 19
-----
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128
-----
Port 20 21 22 23 24 25 26 27 28 29
-----
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128
-----
Port 30 31 32 33 34 35 36 37 38 39
-----
FID 128 | 128 | 128 | 2 | 2 | 2 | 128 | 128 | 128 | 128
```

To display the logical switch configuration for GbE ports only (in the example, all GbE ports are in logical switch 2):

```
switch:admin> lscfg --show -ge

Created switches FIDs(Domain IDs): (ds) 2(bs)(1) 1(1)

Slot 1 2 3 4 5 6 7 8
-----
Port
0 | | 2 | | | | | | |
1 | | 2 | | | | | | |
2 | | 2 | | | | | | |
3 | | 2 | | | | | | |
4 | | 2 | | | | | | |
5 | | 2 | | | | | | |
6 | | 2 | | | | | | |
```

```

7 | | 2 | | | | | | | |
8 | | 2 | | | | | | | |
9 | | 2 | | | | | | | |
10 | | 2 | | | | | | | |
11 | | 2 | | | | | | | |

```

To display the partition configuration for all slots with the **-ge** option:

```

switch:admin> lscfg --show -provision -ge

Created switches FIDs(Domain IDs): 128(ds)(118) 2(bs)(1) 1(1)
Slot      1      2      3      4      5      6      7      8
-----
Port
0 | 128 | 2 | | | | | | 128 | 128 |
1 | 128 | 2 | | | | | | 128 | 128 |
2 | 128 | 2 | | | | | | 128 | 128 |
3 | 128 | 2 | | | | | | 128 | 128 |
4 | 128 | 2 | | | | | | 128 | 128 |
5 | 128 | 2 | | | | | | 128 | 128 |
6 | 128 | 2 | | | | | | 128 | 128 |
7 | 128 | 2 | | | | | | 128 | 128 |
8 | 128 | 2 | | | | | | 128 | 128 |
9 | 128 | 2 | | | | | | 128 | 128 |
10 | 128 | 2 | | | | | | 128 | 128 |
11 | 128 | 2 | | | | | | 128 | 128 |

```

To change the fabric ID for a logical switch:

```

switch:admin> lscfg --change 1-newfid 2
Changing of a switch fid requires that the switch be disabled.
Would you like to continue [y/n]?: y
Disabling switch...
All active login sessions for FID 2 have been terminated.
Checking and logging message: fid = 2.
Please enable your switch.

```

To display the change:

```

switch:admin> lscfg --show

Created switches FIDs(Domain IDs): 128(ds)(118) 1(1) 2(bs)(1)

Port  0      1      2      3      4      5      6      7      8      9
-----
FID   1 | 1 | 1 | 1 | 1 | 128 | 128 | 128 | 128 | 128 |
Port 10     11     12     13     14     15     16     17     18     19
-----
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |
Port 20     21     22     23     24     25     26     27     28     29
-----
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |
Port 30     31     32     33     34     35     36     37     38     39
-----
FID 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 | 128 |

```

To make logical switch FID 1 the base switch without confirmation:

```
switch:admin> lscfg --change 1 -base -force
Disabling the current base switch...
Disabling switch fid 1
Disabling the proposed new base switch...
Disabling switch fid 1
Please enable your switches when ready.
```

To make logical switch FID 1 the base switch with confirmation:

```
switch:admin> lscfg --change 1 -base
Creation of a base switch requires that the proposed new base switch on this
system be disabled.

Would you like to continue [y/n]?: y
Disabling the proposed new base switch...
Disabling switch fid 1
Please enable your switches when ready.
```

To display the logical switch details with the FID name:

```
switch:admin> lscfg --show -n
-----
Switch Information
-----
FID: 30
SwitchType: DS
DomainID: 3
SwitchName: Pluto2
FabricName: thor2fab
-----
FID: 2
SwitchType: BS
DomainID: 1
SwitchName: switch_2
FabricName: base_switch
```

SEE ALSO [IfCfg, setContext](#)

IsDbShow

Displays the Fabric Shortest Path First (FSPF) link state database.

Synopsis `Isdbshow [domain]`

Description Use this command to display an FSPF link state database record for switches in the fabric or for a specified domain.

There are two types of database entries:

- The link state database entry, which is permanently allocated.
- The link state record (LSR), which is allocated when a switch is connected to the fabric.

The LSR describes the links between connected domains in a fabric. For a link to be reported in the LSR, the neighbor for that link must be in NB_ST_FULL state.

This command displays the content of both types of database entries, if both are present, as shown below:

| | |
|----------------------|--|
| Domain | Domain ID described by this LSR. A (self) keyword after the domain ID indicates that LSR describes the local switch. |
| IsrP | Pointer to LSR. |
| earlyAccLSRs | Number of LSRs accepted, even though they were not sufficiently spaced apart. |
| ignoredLSRs | Number of LSRs not accepted because they were not sufficiently spaced apart. |
| lastIgnored | Last time an LSR was ignored. |
| installTime | Time this LSR was installed in the database, in seconds since boot. |
| lseFlags | Internal variable. |
| uOutlfsP | Internal variable. |
| uAllOutlfsP | Internal variable. |
| uPathCost | Internal variable. |
| uOldHopCount | Internal variable. |
| uHopsFromRoot | Internal variable. |

| | |
|----------------------------------|---|
| uPathCount | The number of currently available paths to the remote domain. |
| mOutIfsP | Internal variable. |
| parent | Internal variable. |
| mPathCost | Internal variable. |
| mHopsFromRoot | Internal variable. |
| Link State Record pointer | Pointer to LSR. The same as lsrP. |
| lsAge | Age, in seconds, of this LSR. An LSR is removed from the database when its age exceeds 3,600 seconds. |
| reserved | Reserved for future use. |
| type | Type of the LSR. Always 1. |
| options | Always 0. |
| lsId | ID of this LSR. It is identical to the domain ID. |
| advertiser | Domain ID of the switch that originated this LSR. |
| incarn | Incarnation number of this LSR. |
| length | Total length, in bytes, of this LSR. Includes header and link state information for all links. |
| chksum | Checksum of total LSR, with exception of lsAge field. |
| linkCnt | Number of links in this LSR. Each link represents a neighbor in NB_ST_FULL state. |
| flags | Always 0. |
| LinkId | ID of this link. It is the domain ID of the switch on the other side of the link. |
| out port | Port number on the local switch. |

| | |
|-----------------|---|
| rem port | Port number of the port on the other side of the link. |
| cost | Cost of this link. The default cost for a 1 Gbps link is 1,000. |
| bw | The rounded bandwidth of the output link, in Gbps. |
| type | Always 1. |

Notes Beginning Fabric OS v7.3.0, the output displays only the lines with a bit set. If a port bitmap does not have any bits set, the output displays as "None" for the first line of the bitmap.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

| | |
|---------------|--|
| domain | Specifies the domain ID of the LSR to be displayed. This operand is optional; if omitted, the entire link state database is displayed. |
|---------------|--|

Examples To display the link state record for a switch:

```
switch:admin> lsdbshow 1

Domain = 1 (self), Link State Database Entry pointer = 0x1096da60
lsrP           = 0x109784b0
earlyAccLSRs  = 0
ignoredLSRs   = 0
lastIgnored   = Never
installTime   = Aug 26 18:20:41.451
lseFlags      = 0xa
uAllOutIfsP[4] = 0x00000001
uPathCost     = 0
uOldHopCount  = 0
uHopsFromRoot = 0
uPathCount    = 1
mOutIfsP[0]   = 0x10000000
parent        = 0xf0
mPathCost     = 0
mHopsFromRoot = 0

Link State Record:
Link State Record pointer = 0x109784b0
lsAge         = 321
reserved      = 0
type          = 1
options       = 0x0
lsId         = 1
advertiser    = 1
incarn       = 0x80000185
```

```
length          = 60
chksum          = 0x168a
linkCnt = 2,    flags = 0x0
LinkId = 91, out port = 28, rem port = 28, cost = 500, bw = 16G, type = 1
LinkId = 91, out port = 29, rem port = 29, cost = 500, bw = 48G, type = 1
```

SEE ALSO [interfaceShow](#), [nbrStateShow](#)

mapsConfig

Sets MAPS configurations.

Synopsis

```
mapsconfig --config pause -type member_type
           -members member_list

mapsconfig --config continue -type member_type
           -members member_list

mapsconfig --actions actions_list

mapsconfig --emailcfg -address email_address

mapsconfig --testmail [-subject subject] [-message msg]

mapsconfig --import flow_name [-force]

mapsconfig --deimport flow_name [-force]

mapsconfig --purge

mapsconfig --show

mapsconfig --help
```

Description Use this command to perform the following Monitoring and Alerting Policy Suite (MAPS) functions:

- Pause or continue monitoring specific elements.
- Define the list of allowable actions that can be taken on the switch when a threshold is triggered.
- Configure e-mail address to which the alerts must be delivered.
- Import a flow.
- Remove the imported flow from MAPS configuration.
- Delete all user-defined MAPS configurations related to rules, groups, policies, and so on.
- Display MAPS settings.

Notes This command requires a Fabric Vision license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

```
--config pause           Stops monitoring specific elements. The active policy is still in effect.

--config continue       Resumes monitoring specific elements on which monitoring was previously paused.
```

- type *member_type***
Specifies the type of elements on which to pause monitoring. Valid types include the following:
- port**
The elements are FC ports.
 - circuit**
The elements are FCIP circuits.
 - sfp**
The elements are SFP transceivers.
- members *member_list***
Specifies the specific elements on which to pause or continue monitoring. The member list must be enclosed in double quotation marks and can consist of the following:
- Port and SFP groups:
- A single port, for example, "8" or "2/8" on bladed systems.
 - A port range where the beginning and end port are separated by a dash, for example, "8-13" or "2/8-13" on bladed systems.
 - A set of ports and port ranges separated by commas, for example, "3,5,7-10,15" or "2/3,2/5,3/7-10,4/15" on bladed systems.
- Circuit groups:
- A single circuit in VE_Port/circuit format, for example, "33/5".
 - A circuit range in VE_Port/circuit format, for example, "33/1-3".
 - A set of circuits separated by commas, for example, "33/1-3,33/5".
- actions *actions_list***
Specifies a comma-separated list of all of the actions that are allowed globally on the switch. Any actions listed for a specific rule must also be listed here to take effect. Valid actions include the following:
- raslog**
Generates a RASLog message.
 - decom**
Decommissions the port. Use this action along with **fence**.
 - fence**
Fences the port, if port fencing is enabled. Port fencing takes the ports offline if the rule thresholds are exceeded. This action is valid only on conditions that are evaluated on ports.
 - snmp**
Generates an SNMP trap.
 - email**
Sends information about a switch event to a specified e-mail address.
 - sfp_marginal**
Sets the state of the affected SFP transceiver to marginal. This action is valid only if the monitoring system (specified by the monitor operand in the **mapsRule** command) is one of the switch policy status monitoring systems: CURRENT, RXP, TXP, VOLTAGE, SFP_TEMP.

- sddq**
Isolates the slow-drain flows to a low priority VC from the existing VC (medium or high) thus freeing up the resources for the regular flows in the existing VC. Th slow drain device quarantine (SDDQ) feature is not supported in the Access Gateway mode.
- fms**
Notifies the configured MAPS threshold events to the FICON Management Server (FMS).
- toggle**
Enables port toggling to recover a port from bottleneck condition caused by the target device. Port toggle is supported only for the F_Ports.
- none**
No actions are allowed on the switch. Specifying this option allows you to turn off all notifications.

The actions SW_CRITICAL and SW_MARGINAL are always enabled by default.
- emailcfg**
Specifies the configuration for e-mail notifications, if **email** is a triggered action. You can configure a relay host IP address using the **relayConfig** command.
- address *email_address***
Specifies the e-mail address to which the notifications are sent. You can configure up to 5 e-mail addresses as the alert message's receiver. Multiple e-mail addresses must be separated by a comma and each e-mail address can be up to a maximum length of 128 bytes. The e-mail address must conform to standard syntax: string@domain.suffix. Invalid e-mail addresses are rejected. "NONE" is the default address and a valid input parameter.
- testmail**
Sends the test e-mail with the default subject and message along with the switch name in the message. The command fails if the e-mail address is not configured.
- subject *subject***
Specifies the user-defined subject line for the test e-mail. This operand is optional.
- message *msg***
Specifies the user-defined message for the test e-mail. This operand is optional.
- import**
Imports a flow from Flow Vision into MAPS.
- flow_name***
Specifies the name of a flow to be imported. The flow name must be defined in Flow Vision. This operand is required.
- force**
Imports a flow with the same name as the previously imported flow. This operand is optional. When a flow that was imported into MAPS is deleted in Flow Vision, all members of the group corresponding to that flow are automatically deleted from MAPS. If the same flow is created again, it will not automatically be reimported into MAPS, and importing this flow using the **--import** option fails unless you use the **-force** option.

--deimport

Removes a flow from MAPS. The following operand is required:

***flow_name* [-force]**

Specifies the name of the flow to be removed from MAPS. The **-force** operand is optional. You cannot remove a flow that is used by any rules. The **-force** option overrides the default behavior. If a flow is present in user-defined rules, the **-force** option deletes all the rules that are configured with the given flow and then removes the flow.

--purge

Deletes all user-defined MAPS configurations on the switch (groups, rules, and policies). This command does not delete the predefined groups and policies. Enables **dflt_conservative_policy** after successful purge.

--show

Displays the MAPS global configuration settings.

--help

Displays the command usage.

Examples

To stop monitoring on three ports:

```
switch:admin> mapsconfig --config pause -type port
                -members "3/1,3/2,3/3"
```

To resume monitoring on two ports:

```
switch:admin> mapsconfig --config continue -type port
                -members "3/1,3/3"
```

To specify that generating a RASLog and sending an e-mail message are the only allowed notification actions on this switch:

```
switch:admin> mapsconfig --actions raslog,email
```

To configure the e-mail address to which notifications are sent:

```
switch:admin> mapsconfig --emailcfg -address admin@mycompany.com
```

To send a test e-mail:

```
switch:admin> mapsconfig --testmail
MAPS test welcome mail sent successfully
```

To import a flow:

```
switch:admin> mapsconfig --import fmFlow1
```

To delete all of the user-defined groups, policies, and rules:

```
switch:admin> mapsconfig --purge
WARNING: This command will clear all the user-defined MAPS \
         configurations and activate the factory defined policy \
         and rules.
Do you want to continue? (yes, y, no, n): [no] yes
```

mapsConfig

To display the global MAPS configuration:

```
switch:admin> mapsconfig --show  
Configured Notifications: RASLOG,EMAIL,FENCE,SW_CRITICAL,SW_MARGINAL,SDDQ  
Mail Recipient:         admin@myCompany.com  
=====
```

| | |
|-----------|------------|
| PORT : | 3/25,8/5-9 |
| CIRCUIT : | |
| SFP : | |

SEE ALSO [logicalGroup](#), [mapsDb](#), [mapsPolicy](#), [mapsRule](#), [mapsSam](#), [relayConfig](#)

mapsDb

Displays or clears the dashboard showing an at-a-glance snapshot of switch health status.

Synopsis `mapsdb --show [all | history | details [-day mm/dd/yyyy | -hr hour_of_day]]`

`mapsdb --clear [all | history | summary]`

`mapsdb --help`

Description Use this command to view the summary of the events or rules triggered and the objects on which the rules were triggered over a specified period of time, and to clear the dashboard data. When used with the **--show** option, this command displays the following information:

Dashboard Information

Displays the dashboard start time, active policies, configured alerts, fenced ports, decommissioned ports, fenced circuits, quarantined ports, and the ports with highest zoned device ratio.

Switch Health Report

Displays the overall status of the switch. If the overall status is not healthy, the contributing factors and the fenced ports are listed.

Summary Report

Displays the status of the monitoring categories needed to determine the current health of the switch. The health state of each category is contributed to by a group of monitoring systems or error counters, and the rules configured in the active policy. The summary view displays the following information for each category: information collected since midnight of the current day and the historical information collected over the last 7 days.

The health state can be one of the following:

No Error

No error has occurred.

In operating range

The errors are within the thresholds configured in the active policy.

Out of operating range

The errors are above the configured thresholds and this triggers the rules configured in the active policy. This indicates that some attributes of the switch are operating out of the configured range.

Rules Affecting Health

Displays the conditions that contributed to the current switch status. The following information is displayed: number of rules triggered in a category, repeat count, triggered rules, execution time, the elements for which rules were triggered, and the triggered value.

History Data

Displays the following historical data for a specific time window: the monitoring system, the current value that triggered the rule, and the elements for which rules were triggered.

History Data for Backend ports

Displays the error statistics for the backend ports for a specific time window.

If no operands are specified, this command displays the usage.

Notes

This command requires a Fabric Vision license. Without Fabric Vision license, this command displays the summary of only the unlicensed features.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

--show

Displays a summary of the data collected since midnight of the current day. The following operands are supported with the **--show** command.

all

Displays a summary and historical data of the errors for 5 rules and last 5 ports.

history

Displays the historical data only.

details

Displays detailed data in addition to the summary. The following operands are supported with the **details** command.

-day mm/dd/yyyy

Displays the dashboard data for a specific day.

mm

Specifies the month. Valid values are 01 through 12.

dd

Specifies the date. Valid values are 01 through 31.

yyyy

Specifies the year.

-hr hour_of_day

Displays the dashboard data for a specific hour of the day. The valid values for hour are 0 through 23. For example, if you specify 17, the dashboard data is displayed for the time window of 5:00 PM to 5:59 PM.

| | |
|----------------|---|
| --clear | Clears the dashboard data. This command clears all database data except the History Data of the current day because it is directly displayed from hardware. The following operands are supported with the --clear command. |
| all | Clears all dashboard data. |
| history | Clears the dashboard history data. |
| summary | Clears the dashboard summary data. |
| --help | Displays the command usage. |

Examples To display the dashboard with summary information:

```
switch:admin> mapsdb --show all
1 Dashboard Information:
=====

DB start time:           Mon Jan  5 08:36:08 2015
Active policy:          dflt_aggressive_policy
Configured Notifications: RASLOG,SW_CRITICAL,SW_MARGINAL,SDDQ
Fenced Ports :         None
Decommissioned Ports : None
Fenced circuits :      38/0,38/1,38/2,38/3,38/4,38/5
Quarantined Ports :    8/0,8/4
Ports with highest Zoned device ratio: 1,2,3

2 Switch Health Report:
=====

Current Switch Policy Status: HEALTHY

3.1 Summary Report:
=====

Category                |Today                |Last 7 days          |
-----|-----|-----|
Port Health              |No Errors            |In operating range   |
BE Port Health           |No Errors            |No Errors            |
Fru Health               |In operating range   |In operating range   |
Security Violations      |No Errors            |No Errors            |
Fabric State Changes     |No Errors            |In operating range   |
Switch Resource          |In operating range   |In operating range   |
Traffic Performance      |In operating range   |In operating range   |
FCIP Health              |Not applicable       |Not applicable       |
Fabric Performance Impact|In operating range   |In operating range   |
Conditions contributing to health:
3.2 Rules Affecting Health:
=====

Category(Rule Count)|RepeatCount|Rule Name |Execution Time |Object |Triggered
```

```

Value(Units) |
-----
-----

4 History Data:
=====

Stats(Units)      Current    01/05/15   --/--/--   --/--/--   --/--/--
                  Port(val)  Port(val)
-----
CRC(CRCs)         -          -          -          -          -
ITW(ITWs)         -          10(95)     -          -          -
LOSS_SYNC(SyncLoss) -          -          -          -          -
LF                -          10(4)     -          -          -
LOSS_SIGNAL(LOS)  -          10(48)    -          -          -
PE(Errors)        -          -          -          -          -
STATE_CHG         -          10(48)    -          -          -
LR                -          10(32)    -          -          -
C3TXTO(Timeouts) -          -          -          -          -
RX(%)             -          -          -          -          -
TX(%)             -          -          -          -          -
UTIL(%)           -          -          -          -          -
CRED_ZERO(Million) -          10(4)     -          -          -
BN_SECS(Seconds) -          -          -          -          -

```

```

5 History Data for Backend ports:
=====

Stats(Units)      Current    01/05/15   --/--/--   --/--/--   --/--/--
-----
CRC(CRCs)         -          -          -          -          -
ITW(ITWs)         -          -          -          -          -
LR                -          -          -          -          -
BAD_OS(Errors)    -          -          -          -          -
FRM_LONG(Errors)  -          -          -          -          -
FRM_TRUNC(Errors) -          -          -          -          -

```

To display the dashboard with summary information and historical data for a specific day:

```

switch:admin> mapsdb --show details -day 01/06/2015
1 Dashboard Information:
=====

DB start time:           Mon Jan  5 08:36:08 2015
Active policy:           dflt_aggressive_policy
Configured Notifications: RASLOG,SW_CRITICAL,SW_MARGINAL,SDDQ
Fenced Ports :          None
Decommissioned Ports :  None
Fenced circuits :       38/0,38/1,38/2,38/3,38/4,38/5
Quarantined Ports :     8/0,8/4
Ports with highest Zoned device ratio: 1,2,3,4

2 Switch Health Report:
=====

Current Switch Policy Status: HEALTHY

3.1 Summary Report:
=====

```

| Category | Today | Last 7 days |
|---------------------------|--------------------|--------------------|
| Port Health | No Errors | In operating range |
| BE Port Health | No Errors | No Errors |
| Fru Health | In operating range | In operating range |
| Security Violations | No Errors | No Errors |
| Fabric State Changes | No Errors | In operating range |
| Switch Resource | In operating range | In operating range |
| Traffic Performance | In operating range | In operating range |
| FCIP Health | Not applicable | Not applicable |
| Fabric Performance Impact | In operating range | In operating range |

Conditions contributing to health:
3.2 Rules Affecting Health:
=====

| Category(Rule Count) | RepeatCount | Rule Name | Execution Time | Object | Triggered |
|----------------------|-------------|-----------|----------------|--------|-----------|
| Value(Units) | | | | | |

4 History Data:
=====

| Stats(Units) | Current | 01/05/15 | --/--/-- | --/--/-- | --/--/-- |
|---------------------|-----------|-----------|----------|----------|----------|
| Port(val) | Port(val) | Port(val) | | | |
| CRC(CRCs) | - | - | - | - | - |
| ITW(ITWs) | - | 10(95) | - | - | - |
| LOSS_SYNC(SyncLoss) | - | - | - | - | - |
| LF | - | 10(4) | - | - | - |
| LOSS_SIGNAL(LOS) | - | 10(48) | - | - | - |
| PE(Errors) | - | - | - | - | - |
| STATE_CHG | - | 10(48) | - | - | - |
| LR | - | 10(32) | - | - | - |
| C3TXTO(Timeouts) | - | - | - | - | - |
| RX(%) | - | - | - | - | - |
| TX(%) | - | - | - | - | - |
| UTIL(%) | - | - | - | - | - |
| CRED_ZERO(Million) | - | 10(4) | - | - | - |
| BN_SECS(Seconds) | - | - | - | - | - |

5 History Data for Backend ports:
=====

| Stats(Units) | Current | 01/05/15 | --/--/-- | --/--/-- | --/--/-- |
|-------------------|---------|----------|----------|----------|----------|
| CRC(CRCs) | - | - | - | - | - |
| ITW(ITWs) | - | - | - | - | - |
| LR | - | - | - | - | - |
| BAD_OS(Errors) | - | - | - | - | - |
| FRM_LONG(Errors) | - | - | - | - | - |
| FRM_TRUNC(Errors) | - | - | - | - | - |

To display the dashboard information without Fabric Vision license:

```
switch:admin> mapsdb --show
1 Dashboard Information:
=====
```

```
DB start time:      Wed Oct  1 06:04:33 2014
Active policy:     dflt_base_policy
```

```
Configured Notifications: SW_CRITICAL,SW_MARGINAL
```

```
2 Switch Health Report:
=====
```

```
Current Switch Policy Status: MARGINAL
```

```
Contributing Factors:
```

```
-----
*FAULTY_BLADE (MARGINAL).
*HA_SYNC (MARGINAL).
```

```
3.1 Summary Report:
=====
```

| Category | Today | Last 7 days |
|-----------------|------------------------|-------------|
| Fru Health | Out of operating range | No Errors |
| Switch Resource | No Errors | No Errors |

```
3.2 Rules Affecting Health:
=====
```

| Category(Rule Count) | RepeatCount | Rule Name |
|----------------------|-------------|----------------------------|
| Fru Health(1) | 1 | defALL_SLOTSBLADE_STATE_IN |

| Execution Time | Object | Triggered Value(Units) |
|-------------------|---------|------------------------|
| 10/01/14 06:09:42 | Blade 8 | IN |

MAPS is not Licensed. MAPS extended features are available ONLY with License

To clear the history data:

```
switch:admin> mapsdb --clear history
```

SEE ALSO [logicalGroup](#), [mapsConfig](#), [mapsPolicy](#), [mapsRule](#), [mapsSam](#), [portStatsClear](#)

mapsHelp

Displays MAPS command information.

Synopsis `mapshelp`

Description Use this command to display a listing of Monitoring and Alerting Policy Suite (MAPS) commands with short descriptions for each command. MAPS commands require a Fabric Vision license.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display MAPS command help information:

```
switch:admin> mapshelp
mapsRule          Use this command to create threshold monitoring rules
mapsPolicy        Manage policies using this command
logicalGroup      Create and manage user-defined logical groups
mapsdb            Display MAPS Dashboard to see the summary of the \
                  violations triggered due to current active policy.
mapsConfig        MAPS configuration command eg: use this to enable MAPS
relayConfig       Configure e-mail server information to receive email \
                  notifications
mapssam           Show availability monitor information
mapsHelp          Display all the MAPS commands
```

SEE ALSO None

mapsPolicy

Creates and manages MAPS policies.

Synopsis

```
mapspolicy --create policy_name
mapspolicy --addrule policy_name -rulename rule_name
mapspolicy --delrule policy_name -rulename rule_name
mapspolicy --enable policy_name
mapspolicy --clone existing_policy_name -name new_policy_name
mapspolicy --delete policy_name
mapspolicy --show policy_name | -summary | -all [-concise]
mapspolicy --help
```

Description Use this command to create and manage monitoring policies. A MAPS policy is a set of rules that define thresholds for measures and actions to take when a threshold is triggered. When you enable a policy, all of the rules in the policy are in effect. A switch can have multiple policies.

Notes This command requires a Fabric Vision license.

Any change to a policy affects only the information stored on the switch. Consequentially, if an active policy is modified, it has to be re-enabled for the modifications to be effective.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--create Creates a monitoring policy.

policy_name Specifies the name of the policy to be created. The name for the policy must be unique; it is case-sensitive and can contain up to 32 characters.

--addrule *policy_name* Adds a rule to the specified user-defined policy. You cannot add a rule to the preconfigured policies.

-rulename *rule_name* Specifies the name of the rule to be added to the policy.

--delrule *policy_name* Deletes a rule from the specified user-defined policy. You cannot delete a rule from the preconfigured policies.

-rulename *rule_name* Specifies the name of the rule to be deleted from the policy.

- enable**
Activates a policy on the switch. The dashboard data will be cleared after the policy is activated.
- policy_name*
Specifies the name of the policy to be enabled. The name can be a predefined or a user-defined policy. If another policy is already enabled, that policy is disabled and the specified policy is enabled instead. The predefined policies are as follows:
- dflt_conservative_policy**
Contains rules with more lenient thresholds that allow a buffer and do not immediately trigger actions. Use this policy in environments where the elements are resilient and can accommodate errors.
- dflt_aggressive_policy**
Contains rules with very strict thresholds. Use this policy if you need a pristine fabric (for example, FICON fabrics).
- dflt_moderate_policy**
Contains rules with thresholds values between the aggressive and conservative policies.
- dflt_base_policy**
Contains rules based on the features which can be monitored without a license.
- clone**
Creates a replica of an existing predefined or user-defined policy. The new policy has all of the rules of the existing policy. You can further modify the newly created policy. The following operands are required:
- existing_policy_name*
Specifies the name of an existing policy. The policy can be a predefined policy or a user-defined policy.
- name new_policy_name**
Specifies the name of the policy to be created.
- delete**
Deletes a user-defined policy.
- policy_name*
Specifies the name of the user-defined policy to be deleted. You cannot delete the predefined policies.
- show**
Displays the rules in a policy.
- policy_name*
Specifies the name of the policy to display. The name can be a predefined or a user-defined policy.
- summary**
Displays all existing policies and the number of rules present in them.
- all**
Displays the rules in all of the policies.

-concise

Displays legends instead of the complete action names in the output. The legends for the action names are as follows: RS (RASLOG), PF (FENCE), EML (EMAIL), SN (SNMP), PT (TOGGLE), PD (DECOM), FMS (FMS), SDDQ (SDDQ), SWD (SW_CRITICAL), SWM (SW_MARGINAL), and SFPM (SFP_MARGINAL). This operand is optional.

--help

Displays the command usage.

Examples

To create a policy named "aggressive_policy":

```
switch:admin> mapspolicy --create aggressive_policy
```

To add a rule to a policy:

```
switch:admin> mapspolicy --addrule aggressive_policy \
  -rulename CRIT_PORT_CRC_GE_2
```

To delete a rule from a policy:

```
switch:admin> mapspolicy --delrule aggressive_policy \
  -rulename CRIT_PORT_CRC_GE_2
```

To clone a policy:

```
switch:admin> mapspolicy --clone dflt_conservative_policy \
  -name my_conservative_policy
```

To delete a policy:

```
switch:admin> mapspolicy --delete my_conservative_policy
```

To display a policy:

```
switch:admin> mapspolicy --show Policy1
Policy Names  Rule Names  Action  Condition
-----
Policy1      Rule1      RASLOG  Switch(SEC_IDB/Min>0)
              Rule2      RASLOG  Switch(SEC_IDB/Min>1)
Active Policy: Policy2
```

To display a summary of policies:

```
switch:admin> mapspolicy --show -summary
Policy Names      Number of Rules
-----
Policy1          :          125
Policy2          :          130
Policy3          :          135
Policy4          :          150
Active Policy: Policy2
```

To enable a policy:

```
switch:admin> mapspolicy --enable aggressive_policy
```

To display the details of a policy:

```
switch:admin> mapspolicy --show dflt_base_policy -concise
Policy Name: dflt_base_policy

Rule Name                                     |Condition                                     |Actions
-----|-----|-----
-----|-----|-----
defALL_TSTEMP_OUT_OF_RANGE                   | ALL_TS (TEMP/NONE==OUT_OF_RANGE)
| RS, SN, EML |
defCHASSISFLASH_USAGE_90                     | CHASSIS (FLASH_USAGE/NONE>=90)
| RS, SN, EML |
defCHASSISMEMORY_USAGE_75                   | CHASSIS (MEMORY_USAGE/NONE>=75)
| RS, SN, EML |
defCHASSISCPU_80                             | CHASSIS (CPU/NONE>=80)
| RS, SN, EML |
defCHASSISBAD_TEMP_MARG                     | CHASSIS (BAD_TEMP/NONE>=1)
| SWM, SN, EML |
defCHASSISBAD_TEMP_CRIT                     | CHASSIS (BAD_TEMP/NONE>=2)
| SWD, SN, EML |
defCHASSISBAD_PWR_CRIT                      | CHASSIS (BAD_PWR/NONE>=2)
| SWD, SN, EML |
defCHASSISBAD_FAN_MARG                      | CHASSIS (BAD_FAN/NONE>=1)
| SWM, SN, EML |
defCHASSISBAD_FAN_CRIT                      | CHASSIS (BAD_FAN/NONE>=2)
| SWD, SN, EML |
defALL_PSPS_STATE_FAULTY                    | ALL_PS (PS_STATE/NONE==FAULTY)
| RS, SN, EML |
defALL_PSPS_STATE_ON                        | ALL_PS (PS_STATE/NONE==ON)
| RS, SN, EML |
defALL_PSPS_STATE_OUT                       | ALL_PS (PS_STATE/NONE==OUT)
| RS, SN, EML |
defALL_FANFAN_STATE_FAULTY                  | ALL_FAN (FAN_STATE/NONE==FAULTY)
| RS, SN, EML |
defALL_FANFAN_STATE_ON                      | ALL_FAN (FAN_STATE/NONE==ON)
| RS, SN, EML |
defALL_FANFAN_STATE_OUT                     | ALL_FAN (FAN_STATE/NONE==OUT)
| RS, SN, EML |
*defALL_PORTSSFP_STATE_FAULTY               | ALL_PORTS (SFP_STATE/NONE==FAULTY)
| RS, SN, EML |
*defALL_PORTSSFP_STATE_OUT                  | ALL_PORTS (SFP_STATE/NONE==OUT)
| RS, SN, EML |
*defALL_PORTSSFP_STATE_IN                   | ALL_PORTS (SFP_STATE/NONE==IN)
| RS, SN, EML |
defCHASSISETH_MGMT_PORT_STATE_DOWN          | CHASSIS (ETH_MGMT_PORT_STATE/NONE==DOWN)
| RS, SN, EML |
defCHASSISETH_MGMT_PORT_STATE_UP           | CHASSIS (ETH_MGMT_PORT_STATE/NONE==UP)
| RS, SN, EML |

Active Policy is 'dflt_base_policy'.
Unmonitored Rules are prefixed with "*"

Legends:
RS: RASLOG SN: SNMP EML: EMAIL PF: FENCE PL: PORTLOG PD: DECOM FMS: FMS PT: TOGGLE
SDDQ: SDDQ SWD: SW_CRITICAL
```

SEE ALSO [logicalGroup](#), [mapsConfig](#), [mapsDb](#), [mapsRule](#), [mapsSam](#)

mapsRule

Creates and manages monitoring rules.

Synopsis

```
mapsrule --create rule_name -group group_name | flow_name
  -monitor ms_name -timebase timebase
  -op comparison_operator -value value
  -action action -tt timeout
  -qt quiet_time [-policy policy]

mapsrule --delete rule_name [-force]

mapsrule --config rule_name [-group group_name | flow_name]
  [-monitor monitor] [-timebase timebase]
  [-op comparison_operator] [-value value] [-action action]
  -tt timeout -qt quiet_time | -qtclear

mapsrule --clone existing_rule_name -rulename new_rule_name
  [-group group_name | flow_name] [-monitor ms_name]
  [-timebase timebase] [-op comparison_operator]
  [-value value] [-action action] -tt timeout
  -qt quiet_time | -qtclear [-policy policy]

mapsrule --show rule_name | -all [-concise]

mapsrule --help
```

Description Use this command to configure and manage MAPS monitoring rules, and to display the configured rules. A rule associates a condition with actions that must be triggered when the specified condition is evaluated to be true. When you modify a rule, the rule does not take effect until you enable the policy. If the rule is part of the enabled policy, you must re-enable the policy for the rule to take effect.

For a monitoring rule created with an imported flow group, MAPS monitors the paused members.

Notes This command requires a Fabric Vision license.

The port fence and port decommission actions are not supported on an E_Port for a C3TXTO monitoring system.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

rule_name

Specifies the name of the rule to be configured. The name for the rule must be unique; it is case-sensitive and can contain up to 40 characters.

monitor

Specifies the monitoring system. Valid monitoring systems for each category include the following:

Port health

Monitors port statistics. Valid values for the port health category include the following:

| | |
|------------------------|--------------------------------|
| CRC | Cyclic redundancy check errors |
| ITW | Invalid transmit words |
| LOSS_SYNC | Loss of synchronization |
| LF | Link failure |
| LOSS_SIGNAL | Loss of signal |
| PE | Protocol errors |
| LR | Link reset |
| C3TXTO | Class 3 timeouts |
| STATE_CHG | State changes |
| CURRENT | SFP current |
| RXP | SFP receive power |
| TXP | SFP transmit power |
| VOLTAGE | SFP voltage |
| SFP_TEMP | SFP temperature |
| PWR_HRS | SFP power on hours |
| DEV_NPIV_LOGINS | NPIV device logins |

Backend Port Health

Monitors backend port statistics. Valid values for the port health category include the following:

| | |
|--------------------|--|
| CRC | Cyclic redundancy check errors |
| ITW | Invalid transmit words |
| LR | Link reset |
| BAD_OS | Invalid ordered set |
| FRAME_TRUNC | The frame is too short (less than 36 bytes). |
| FRAME_LONG | The frame is longer than expected (greater than 2148 bytes). |

GE Port Health

Monitors backend port statistics. Valid values for the port health category include the following:

| | |
|----------------------|--------------------------------|
| GE_CRC | Cyclic redundancy check errors |
| GE_INV_LEN | Invalid length |
| GE_LOS_OF_SIG | Loss of signal |
| STATE_CHG | Tunnel state change |
| UTIL | Tunnel FC throughput |
| IP_UTIL | Tunnel IP throughput |

FRU health

Monitors the field-replaceable units, including ports, power supplies, and flash memory. Valid values for the FRU health category include the following:

| | |
|--------------------|--------------------------------|
| PS_STATE | Power supply state has changed |
| FAN_STATE | Fan state has changed |
| BLADE_STATE | Blade state has changed |

SFP_STATE SFP state has changed

WWN WWN card state has changed

FCIP health

Monitors the FCIP circuit. Valid values for the FCIP health category include the following:

CIR_STATE Circuit state changes

CIR_UTIL Percentage of circuit utilization

CIR_PKTLOSS Percentage of circuit packet loss

RTT Circuit round trip time in milliseconds

JITTER Percentage of variance in RTT for circuits

STATE_CHG Tunnel state changes

UTIL Percentage of utilization

PKTLOSS Percentage of Tunnel Qos utilization

IP_UTIL Circuit IP utilization

IP_PKTLOSS Circuit IP packet loss

IP_RTT Circuit IP round-trip time in milliseconds

IP_JITTER Circuit IP connection variance

DEV_LATENCY_IMPACT Fabric performance impact

RX Receive bandwidth usage percentage

TX Transmit bandwidth usage percentage

Fabric performance impact

Monitors the fabric performance. Valid values for the Fabric performance impact category include the following:

CRED_ZERO

Buffer credit zero counter increments are indirect indication of latency when frames could not be transmitted on a port due to the delay in receiving R_RDY frames.

THPUT_DEGRADE

Throughput degradation

UTIL

Trunk utilization

IT_FLOW

IT flow ration

Security health

Monitors security violations on the switch. Valid values for the security health category include the following:

SEC_DCC

DCC violations

SEC_HTTP

HTTP violations

SEC_CMD

Illegal command

SEC_IDB

Incompatible security database

SEC_LV

Login violations

SEC_CERT

Invalid certificates

SEC_FCS

No FCS

SEC_SCC

SCC violations

SEC_AUTH_FAIL

Authentication failure

SEC_TELNET

Telnet violations

SEC_TS

Time Server (TS) out of sync

DAYS_TO_EXPIRE

Days to expire

EXPIRED_CERTS
Expired certificates

DID_CHG
Domain ID change

FLOGI
FLOGI

FAB_CFG
Fabric reconfigurations

Fabric state change

Monitors fabric state changes, including zone changes, fabric segmentation, E_Port down, fabric reconfiguration, domain ID changes, and fabric logins. Valid values for the fabric state change category include the following:

DID_CHG
Domain ID change

FLOGI
Fabric logins

FAB_CFG
Fabric reconfigurations

EPORT_DOWN
E_Ports down

FAB_SEG
Fabric segmentation

ZONE_CHG
Zone changes

L2_DEVCNT_PER
Layer 2 device count

LSAN_DEVCNT_PER
LSAN Device Count

ZONE_CFGSZ_PER
Zone configuration size

BB_FCR_CNT
FCR count

TEMP
Temperature sensor

ETH_MGMT_PORT_STATE
Ethernet management port state

FLASH_USAGE
Flash usage

Switch status policy

Monitors the health of the switch. Valid values for the switch policy change category include the following:

FLASH_USAGE

Flash usage

MARG_PORTS

Percentage of marginal ports

FAULTY_PORTS

Percentage of faulty ports

MISSING_SFP

Percentage of missing SFP transceivers

ERR_PORTS

Percentage of error ports

WWN_DOWN

World Wide Name card down

DOWN_CORE

Core blade monitoring

FAULTY_BLADE

Faulty blades

HA_SYNC

High Availability monitoring

FAN_AIRFLOW_MISMATCH

Monitors the air flow direction of the power supply fan FRUs and blower FRUs and generates an alert if there is a mismatch in the air flow direction of any two power supply fans or any two blowers. The mismatch event is followed by a match event, when the air flow direction in all the fans and blowers is back to normal.

Switch resource

Monitors system RAM, flash, memory, and CPU. Valid values for the switch resource category include the following:

BAD_PWR

Absent or faulty power system

BAD_TEMP

Temp Sensors outside range

BAD_FAN

Absent or faulty fans

CPU

CPU usage

MEMORY_USAGE

Memory usage

Traffic performance

Monitors traffic load and flow on an FC Port. Valid values for the traffic performance category include the following:

| | |
|--------------------|------------------------|
| TX_FCNT | Transmit frame count |
| RX_FCNT | Receive frame count |
| TX_THPUT | Transmit throughput |
| RX_THPUT | Receive throughput |
| IO_RD | IO read command count |
| IO_WR | IO write command count |
| IO_RD_BYTES | IO read data |
| IO_WR_BYTES | IO write data |

group_name Specifies the name of a user-defined group.

flow_name Specifies the name of a user-defined flow.

timebase Specifies the time interval between two samples to be compared. Valid time base values include the following:

min Samples are compared every minute.

hour Samples are compared once an hour.

day Samples are compared once a day.

none Time base is not applicable.

comparison_operator Specifies the relational operations to be used in evaluating the condition. Valid operators include the following:

| Less than

| | |
|---------------------|--|
| le | Less than or equal to |
| g | Greater than |
| ge | Greater than or equal to |
| eq | Equal to |
| value | <p>Specifies the threshold value. Valid values include the following:</p> <ul style="list-style-type: none"> • For numerical values: 1-9999. • For percentage values: 1-100. • For FRU states: ON, OFF, IN, OUT, or FAULTY. • For temperature monitoring: IN_RANGE or OUT_OF_RANGE. • For FPI states: IO_FRAME_LOSS, IO_PERF_IMPACT, IO_LATENCY_CLEAR. • For Ethernet port state: UP or DOWN |
| action | <p>Specifies a comma-separated list of actions to be taken if the specified condition is evaluated to be true. The action is not taken unless it is also specified globally, using the mapsconfig --actions command. Valid actions include the following:</p> |
| raslog | Generates a RASLog message. |
| decom | Decommissions the port. Use this action along with fence . |
| fence | Fences the port, if port fencing is enabled. Port fencing takes the ports offline if the user-defined thresholds are exceeded. This action is valid only on conditions that are evaluated on ports. |
| snmp | Generates an SNMP trap. |
| email | Sends information about a switch event to a specified e-mail address. |
| sw_critical | Sets the state of the affected switch to critical. This action is valid only if the monitor operand is one of the switch policy status monitoring systems. |
| sw_marginal | Sets the state of the affected switch to marginal. This action is valid only if the monitor operand is one of the switch policy status monitoring systems. |
| sfp_marginal | Sets the state of the affected SFP transceiver to marginal. This action is valid only if the monitor operand is one of the switch policy status monitoring systems. |

- sddq** Isolates the slow-drain flows to a low priority VC from the existing VC (medium or high) thus freeing up the resources for the regular flows in the existing VC. The slow drain device quarantine (SDDQ) feature is not supported in the Access Gateway mode.
- fms** Notifies the configured MAPS threshold events to the FICON Management Server (FMS).
- toggle** Enables port toggling to recover a port from bottleneck condition caused by the target device. Port toggle is supported only for the F_Ports.
- none** No actions are allowed on the switch. Specifying this option allows you to turn off all notifications.
- create *rule_name*** Creates a monitoring rule and adds it to a specified policy. All the following operands are required:
- group *group_name*** Specifies the group containing the elements on which the specified condition is to be evaluated.
 - monitor *ms_name*** Specifies the monitoring system (MS).
 - timebase *time_base*** Specifies the time interval between two samples to be compared.
 - op *comparison_operator*** Specifies the relational operation to be used in evaluating the condition.
 - value *value*** Specifies the threshold value.
 - action *action*** Specifies a comma-separated list of actions to be taken if the specified condition is evaluated to be true.
 - tt *timeout*** Specifies the time window in seconds over which the port is disabled and re-enabled by MAPS to recover from bottleneck condition. The valid range is from 2 to 3600 seconds. This parameter is valid with only the **toggle** alert option.
 - qt *quiet_time* [-unit min | hour | day]** Specifies the time interval, in units of minutes, hours, or days, between consecutive alerts. This parameter is valid with only the **raslog** and **email** alert options.
 - qtclear** Clears the configured quiet time for the RASLog and e-mail alerts.
 - policy *policy*** Specifies the user-defined policy to which the rule is to be added. You cannot add a rule to a predefined policy.

- delete *rule_name* [-force]**
Deletes a monitoring rule. The **-force** option provides the capability to override the default behavior. This option allows deleting only user-defined rules. The **-force** option fails if a default rule is input as an argument.
- config *rule_name***
Modifies a monitoring rule. You cannot modify a default rule.

Specify only the operands that you want to modify. Refer to **mapsrule --create** for a listing of operands and their descriptions.
- clone *existing_rule_name* -rulename *new_rule_name***
Creates a replica of an existing rule. The following operands are required:
- existing_rule_name***
Specifies the name of the exiting rule to be copied.
- rulename *new_rule_name***
Specifies the name of the new rule. The name for the new rule must be unique; it is case-sensitive and can contain up to 40 characters.

Refer to **mapsrule --create** for a listing of optional operands and their descriptions. Specify only the operands that you want to modify. If no operands are specified, an exact copy is created.
- show**
Displays a rule, including the condition and actions.
- rule_name***
Specifies the name of the rule to display.
- all**
Displays all rules configured on the switch.
- concise**
Displays legends instead of the complete action names in the output. The legends for the action names are as follows: RS (RASLOG), PF (FENCE), EM (EMAIL), SN (SNMP), PT (TOGGLE), PD (DECOM), FM (FMS), SQ (SDDQ), SC (SW_CRITICAL), SM (SW_MARGINAL), and SFPM (SFP_MARGINAL). This operand is optional.
- help**
Displays the command usage.

Examples

To create a rule for monitoring the number of CRC errors every hour on critical ports and generating a RASlog if the number of CRC errors is greater than 10:

```
switch:admin> mapsrule --create check_crc -monitor crc \  
-group CRITICAL_PORTS -timebase hour -op g -value 10  
-action raslog
```

To modify a rule to change the action to generate a RASlog message and fence the port:

```
switch:admin> mapsrule --config check_crc -monitor crc \  
-group CRITICAL_PORTS -timebase hour -op g -value 10  
-action raslog,fence
```

To delete a rule:

```
switch:admin> mapsrule --delete New_Cntr_High
```

To clone a rule with a modified timebase:

```
switch:admin> mapsrule --clone Rule1 -rulename NewRule2
               -timebase Hour
```

To display a single rule:

```
switch:admin> mapsrule --show defRuleITWAboveHighPF
Rule Data:
-----
RuleName: defRuleITWAboveHighPF
Actions: Raslog
Condition: ALL_PORTS(ITW/MIN>25)
Policies Associated: none
```

To display all of the rules configured on a switch:

```
switch:admin> mapsrule --show -all
Rule Name                |Condition                |Actions
-----
defNON_E_F_PORTSCRC_0    |NON_E_F_PORTS(CRC/MIN>0)|
|RASLOG,SNMP,EMAIL      |
defNON_E_F_PORTSCRC_2    |NON_E_F_PORTS(CRC/MIN>2)|
|FENCE,SNMP,EMAIL       |
defNON_E_F_PORTSCRC_10   |NON_E_F_PORTS(CRC/MIN>10)|
|RASLOG,SNMP,EMAIL      |
defNON_E_F_PORTSCRC_20   |NON_E_F_PORTS(CRC/MIN>20)|
|FENCE,SNMP,EMAIL       |
defNON_E_F_PORTSCRC_21   |NON_E_F_PORTS(CRC/MIN>21)|
|RASLOG,SNMP,EMAIL      |
defALL_HOST_PORTSSTATE_CHG_20|ALL_HOST_PORTS(STATE_CHG/MIN>20)|
|FENCE,DECOM,SNMP,EMAIL|
defALL_HOST_PORTSC3TXTO_2|ALL_HOST_PORTS(C3TXTO/MIN>2)|
|RASLOG,SNMP,EMAIL,FMS  |
defALL_10GSWL_SFPVOLTAGE_3600|ALL_10GSWL_SFP(VOLTAGE/NONE>=3600)|
|SFP_MARGINAL,RASLOG,SNMP,EM|
|                                                                    |AIL
|
defALL_10GSWL_SFPRXP_1999|ALL_10GSWL_SFP(RXP/NONE>=1999)|
|SFP_MARGINAL,RASLOG,SNMP,EM|
|                                                                    |AIL
|
defALL_PORTS_IO_PERF_IMPACT|ALL_PORTS(DEV_LATENCY_IMPACT/NONE==IO_PERF_IMPACT)|RASLOG,SNMP,EMAIL,SDDQ,TOGG|
|                                                                    |LE
tt=2                        |
defALL_PORTS_IO_FRAME_LOSS|ALL_PORTS(DEV_LATENCY_IMPACT/NONE==IO_FRAME_LOSS)|
|RASLOG,SNMP,EMAIL,SDDQ,TOGG|
|                                                                    |LE
tt=2                        |
defALL_PORTS_IO_LATENCY_CLEAR|ALL_PORTS(DEV_LATENCY_IMPACT/NONE==IO_LATENCY_CLEA|RASLOG,SNMP,EMAIL|
```

To display all of the rules configured on a switch with legends for action names:

```
switch:admin> mapsrule --show -all -concise
Rule Name          |Condition          |Actions
-----|-----|-----
defNON_E_F_PORTSCRC_0 |NON_E_F_PORTS (CRC/MIN>0)
|RS,SN,EM          |
defNON_E_F_PORTSCRC_2 |NON_E_F_PORTS (CRC/MIN>2)
|PF,SN,EM          |
defNON_E_F_PORTSCRC_10 |NON_E_F_PORTS (CRC/MIN>10)
|RS,SN,EM          |
defNON_E_F_PORTSCRC_20 |NON_E_F_PORTS (CRC/MIN>20)
|PF,SN,EM          |
defNON_E_F_PORTSCRC_21 |NON_E_F_PORTS (CRC/MIN>21)
|RS,SN,EM          |
defALL_HOST_PORTSSTATE_CHG_20 |ALL_HOST_PORTS (STATE_CHG/MIN>20)
|PF,PD,SN,EM      |
defALL_HOST_PORTSC3TXTO_2 |ALL_HOST_PORTS (C3TXTO/MIN>2)
|RS,SN,EM,FM      |
defALL_10GSWL_SFPVOLTAGE_3600 |ALL_10GSWL_SFP (VOLTAGE/NONE>=3600)
|SFPM,RS,SN,EM    |
defALL_10GSWL_SFPRXP_1999 |ALL_10GSWL_SFP (RXP/NONE>=1999)
|SFPM,RS,SN,EM    |
defALL_PORTS_IO_PERF_IMPACT |ALL_PORTS (DEV_LATENCY_IMPACT/NONE==IO_PERF_IMPACT)
|RS,SN,EM,SQ,PT tt=2 |
defALL_PORTS_IO_FRAME_LOSS |ALL_PORTS (DEV_LATENCY_IMPACT/NONE==IO_FRAME_LOSS)
|RS,SN,EM,SQ,PT tt=2 |
defALL_PORTS_IO_LATENCY_CLEAR |ALL_PORTS (DEV_LATENCY_IMPACT/NONE==IO_LATENCY_CLEA
|RS,SN,EM          |

Legend:
RS:RASLOG, EM:EMAIL, PD:DECOM, PF:FENCE, SC:SW_CRITICAL SM:SW_MARGINAL, FM:FMS,
PT:TOGGLE, SQ:SDDQ SN:SNMP
```

SEE ALSO [logicalGroup](#), [mapsConfig](#), [mapsDb](#), [mapsPolicy](#), [mapsSam](#)

mapsSam

Generates or clears reports to display CPU, RAM, and flash memory usage, and the port status for every physical and virtual Fibre Channel port on the switch.

Synopsis `mapssam --show[cpu | memory | flash]`

`mapssam --clear`

`mapssam --help`

Description Use this command to generate the reports to display CPU, RAM, and flash memory usage, and the port status for every physical and virtual Fibre Channel port on the switch. This command provides an option to clear the reports. This report displays uptime and downtime for each port and enables you to check if a particular port is failing more often than the others. When issued with the **--show** option, the report displays with the following information:

| | |
|---------------------------|---|
| Port | Port number on the local switch. |
| Type | Port type, such as the following: E (E_Port), F (F_Port), U (U_Port), D (disable port), T (E Trunk ports), VE (VE_Port), G (G_Port), DP (persistently disable port), DIA (D_Port), VEX (VEX_Port), or TF (F Trunk ports). |
| Total Up Time | Percent of time the port was up. |
| Total Down Time | Percent of time the port was faulty. |
| Down Occurrence | Number of times the port was faulty. |
| Total Offline Time | Percent of time the port was offline. |

Notes This command requires a Fabric Vision license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|---------------|---|
| --show | Displays the MAPS Service Availability Module (SAM) report. |
| cpu | Displays system CPU usage. |
| memory | Displays system memory usage. |
| flash | Displays system flash usage. |

--clear Clears the SAM report.

--help Displays the command usage.

Examples To generate a port availability monitor report:

```
switch:admin> mapssam --show
```

| Port | Type | Total Up Time (Percent) | Total Down Time (Percent) | Down Occurrence (Times) | Total Offline Time (Percent) |
|------|------|----------------------------|------------------------------|----------------------------|---------------------------------|
| 0 | U | 0 | 0 | 0 | 100 |
| 1 | U | 0 | 0 | 0 | 100 |
| 2 | U | 0 | 0 | 0 | 100 |
| 3 | F | 100 | 0 | 0 | 0 |
| 4 | U | 0 | 0 | 0 | 100 |
| 5 | U | 0 | 0 | 0 | 100 |
| 6 | U | 0 | 0 | 0 | 100 |
| 7 | U | 0 | 0 | 0 | 100 |
| 8 | U | 0 | 0 | 0 | 100 |
| 9 | E | 0 | 0 | 0 | 0 |
| 10 | U | 0 | 0 | 0 | 100 |

(output truncated)

To display CPU usage:

```
switch:admin> mapssam --show cpu
Showing Cpu Usage:
CPU Usage : 2.0%
```

To display memory usage:

```
switch:admin> mapssam --show memory
Showing Memory Usage:
Memory Usage : 51.0%
Used Memory : 257207k
Free Memory : 247121k
Total Memory : 504328k
```

To display system flash usage:

```
switch:admin> mapssam --show flash
Showing Flash Usage:
Flash Usage : 42%
```

To clear the report:

```
switch:admin> mapssam --clear
```

SEE ALSO [logicalGroup](#), [mapsConfig](#), [mapsDb](#), [mapsPolicy](#), [mapsRule](#)

memShow

Displays the amounts of free and used memory in a switch.

Synopsis `memshow [-b | -k | -m]`

Description Use this command to display free and used memory in the switch, as well as the shared memory and buffers used by the kernel.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

- b** Specify to display memory usage in bytes.
- k** Specify to display memory usage in kilobytes.
- m** Specify to display memory usage in megabytes. By default, memory usage is displayed in bytes.

Examples To view the memory usage:

```
switch:admin> memshow
      total      used      free  shared  buffers  cached
Mem:  129740800 112562176 17178624      0    139264 30396416
Swap:           0         0         0

switch:admin> memshow -m
      total      used      free  shared  buffers  cached
Mem:      123       107       16      0         0       28
Swap:     0         0         0
```

SEE ALSO [supportSave](#)

motd

motd

Sets the banner on the chassis.

Synopsis `motd --set string`

`motd --show`

Description Use this command to set the banner on the chassis.

The banner is a string of alphanumeric characters. It is displayed before you log in to a switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

--set *string*

Specify a text string of alphanumeric characters to be displayed before login. The string must be enclosed in double quotation marks. Spaces are allowed. The maximum length is 116 characters.

The message displays either before or after the login prompt, depending on the SSH client implementation. Fabric OS does not control when the MOTD displays.

--show

Displays the chassis-wide banner.

--help

Displays the command usage.

Examples To set a chassis-level banner for the switch:

```
switch:admin> motd --set "This is a \  
chassis-level banner. It displays before the login."
```

To display the banner:

```
switch:admin> motd --show  
This is a chassis-level banner. It displays before the login.
```

SEE ALSO [bannerSet](#), [bannerShow](#)

msCapabilityShow

Displays the Management Server (MS) capabilities.

Synopsis `mscopyshow`

Description Use this command to display the supported capabilities of the Management Server for each switch in the fabric. An asterisk displays next to the name of the local switch.

Notes Reliable commit service (RCS) is a fabric-wide capability and is supported only if all the switches in the fabric support the service.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the supported MS capabilities for each switch in the fabric:

```
switch:admin> mscopyshow
      Switch WWN          Capability      Switch Name
=====
10:00:00:60:69:20:15:71  0x0000008f    "switch1" *
10:00:00:60:69:00:30:05  0x0000008f    "switch2"
```

Capability Bit Definitions:

```
Bit 0: Basic Config Service Supported.
Bit 1: Platform Management Service Supported.
Bit 2: Topology Discovery Service Supported.
Bit 3: Unzoned Name Service Supported.
Bit 4: Fabric Zone Service Supported.
Bit 5: Fabric Lock Service Supported.
Bit 6: Time Service Supported.
Bit 7: RSCN Small Payload Supported.
Bit 8: Reliable Commit Service(RCS) Supported.
Bit 9: Access Gateway Registration/Discovery Supported.
Bit 10: Administrative Domains(AD) Supported.
Others: Reserved.
```

SEE ALSO [msConfigure](#), [msPIMgmtActivate](#), [msPIMgmtDeactivate](#), [msTdDisable](#), [msTdEnable](#), [msTdReadConfig](#)

msConfigure

Configures the Management Server (MS) access control list (ACL).

Synopsis `msconfigure`

Description Use this command to configure the MS Access Control List (ACL). The MS allows a Storage Area Network (SAN) management application to retrieve and administer the fabric and Interconnect Elements, such as switches. This application is located at the Fibre Channel well-known address, 0xFFFFFA.

If the MS ACL is empty (default), The MS is available to all systems connected to the fabric. By populating the MS ACL with one or more world wide names (WWNs), you can restrict access to MS to the specified WWNs.

This command is interactive and provides the following choices:

- 0 Done
- 1 Display the access list
- 2 Add member based on its port/node WWN
- 3 Delete member based on its port/node WWN

When changing the MS ACL by adding or deleting WWNs, you are prompted to save the new configuration to nonvolatile storage. The saved MS ACL becomes effective upon reboot.

The MS ACL is implemented on a per-switch basis and should be configured on the switch to which the management application is directly connected.

Notes When an FCS policy is enabled, the MS ACL is not used. In such a case, access to MS is controlled by security by way of the MS_POLICY configuration.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the MS ACL:

```
switch:admin> msconfigure

0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 1

MS Access List consists of (5): {
20:01:00:60:69:00:60:10
20:02:00:60:69:00:60:10
```

```
20:03:00:60:69:00:60:10
20:02:00:60:69:00:60:03
20:02:00:60:69:00:60:15

0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0

done ...
```

SEE ALSO [msCapabilityShow](#), [msPlatShow](#), [msPIClearDB](#), [msPIMgmtActivate](#), [msPIMgmtDeactivate](#), [msTdDisable](#), [msTdEnable](#), [msTdReadConfig](#), [secPolicyShow](#)

msPlatShow

Displays the Management Server (MS) platform database.

Synopsis `msplatshow`

Description Use this command to display information from the MS platform database. This command displays the name of each platform object with the platform type (GATEWAY, HOST_BUS_ADAPTER, and so forth), associated management addresses, and associated node names.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the MS platform database for a fabric:

```
switch:admin> msplatshow
-----
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
Associated Management Addresses:
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
-----
Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 2
Associated Node Names:
10:00:00:60:69:20:15:79
10:00:00:60:69:20:15:75
```

SEE ALSO [msCapabilityShow](#), [msConfigure](#), [msPlatShowDBCB](#), [msPIClearDB](#), [msPIMgmtActivate](#), [msPIMgmtDeactivate](#)

msPlatShowDBC

Displays the Management Server (MS) platform service database control block.

Synopsis `msplatshowdcb`

Description Use this command to display the control block fields associated with the platform database.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the MS platform service database control block:

```
switch:admin> msplatshowdcb  
Domain    Worldwide Name      Retry Count Exchange Status  
-----  
3: 10:00:00:60:69:51:10:e6      0             0x2  
-----
```

```
msPlDBC. peerWwn == 00:00:00:00:00:00:00:00.  
msPlDBC. psPeerWwn == 00:00:00:00:00:00:00:00.  
msPlDBC. replicate == 0.  
msPlDBC. fabMaySeg == 255.  
msPlDBC. enabled == 1.
```

SEE ALSO [msCapabilityShow](#), [msConfigure](#), [msPlatShow](#), [msPIClearDB](#), [msPIMgmtActivate](#), [msPIMgmtDeactivate](#)

msPIClearDB

Clears the Management Server (MS) platform database on all switches in the fabric.

Synopsis `mspicleardb`

Description Use this command to clear the MS platform database in the entire fabric. Because this operation cannot be undone, it should not be performed unless it is intended to resolve a database conflict between two joining fabrics or to establish an entirely new fabric with an empty database.

Notes When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To clear the MS platform database:

```
switch:admin> mspicleardb

MS Platform Service is currently enabled.
This will erase MS Platform Service Database in the\
entire fabric.

Would you like to continue this operation?
(yes, y, no, n): [no] y

Request to MS Platform DB Clear operation in progress...

*Completed clearing MS Platform Service Database!!
```

SEE ALSO [msCapabilityShow](#), [msConfigure](#), [msPlatShow](#), [msPlatShowDBCB](#), [msPIMgmtActivate](#), [msPIMgmtDeactivate](#)

msPIMgmtActivate

Activates the Management Server (MS) platform service.

Synopsis `msplmgmtactivate`

Description Use this command to activate the MS platform service throughout the fabric. This command attempts to activate the MS platform service for each switch in the fabric. The change takes effect immediately and is committed to the configuration database of each affected switch. MS activation is persistent across power cycles and reboots.

Notes By default, the MS platform service is disabled.

Before issuing this command, run **msCapabilityShow** to verify that all switches in the fabric support the MS platform service; if one switch does not support the service, the command fails.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To activate the MS platform service:

```
switch:admin> msplmgmtactivate

Request to activate MS Platform Service in progress.....

*Completed activating MS Platform Service in the fabric!
```

SEE ALSO [msCapabilityShow](#), [msPlatShow](#), [msPIClearDB](#), [msPIMgmtDeactivate](#)

msPIMgmtDeactivate

Deactivates the Management Server (MS) platform service.

Synopsis `msplmgmtdeactivate`

Description Use this command to deactivate the MS platform service throughout the fabric. This command deactivates the MS platform service for each switch in the fabric and commits the change to nonvolatile storage.

Notes When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To deactivate the MS platform service on all switches in the fabric:

```
switch:admin> msplmgmtdeactivate

MS Platform Service is currently enabled.

This will erase MS Platform Service configuration
information as well as database in the entire fabric.

Would you like to continue this operation?
  (yes, y, no, n): [no] y

Request to deactivate MS Platform Service in progress...

*Completed deactivating MS Platform Service in the fabric!
```

SEE ALSO [msCapabilityShow](#), [msConfigure](#), [msPlatShow](#), [msPlatShowDBCB](#), [msPIClearDB](#), [msPIMgmtActivate](#)

msTdDisable

Disables the Management Server (MS) topology discovery service.

Synopsis `mstdisable ["ALL"]`

Description Use this command to disable the management server topology discovery service on a local switch or an entire fabric. This change takes effect immediately and commits to the configuration database for all affected switches. The change is persistent across power cycles and reboots.

Notes Topology Discovery Management requires the attached devices (including attached switches) to support request node identification data (RNID) extended link service (ELS).

When an FCS policy is enabled, and this command is issued with the "ALL" operand, it can be issued only from the primary FCS.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operand is optional:

"ALL"

Disables the MS topology discovery service throughout the entire fabric. This operand must be enclosed in double quotation marks.

Examples To disable the MS topology discovery service on the local switch only:

```
switch:admin> mstdisable
This may erase all NID entries. Are you sure?
(yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service
in progress....
done.
*MS Topology Discovery disabled locally.
```

To disable MS topology discovery on all the switches in the fabric:

```
primaryfcs:admin> mstdisable "ALL"
This may erase all NID entries. Are you sure?
(yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service
in progress....
done.
*MS Topology Discovery disabled locally.
*MS Topology Discovery Disable Operation Complete!!
```

SEE ALSO [msTdEnable](#), [msTdReadConfig](#)

msTdEnable

Enables the Management Server (MS) topology discovery service.

Synopsis `mstdenable ["ALL"]`

Description Use this command to enable the MS topology discovery service on the local switch or throughout the fabric. The change takes effect immediately and commits to the configuration database for all affected switches. The change is persistent across power cycles and reboots.

Notes Topology Discovery Management requires the attached devices (including attached switches) to support request node identification data (RNID) extended link service (ELS).

When an FCS policy is enabled, and this command is issued with the "ALL" operand, it can be issued only from the primary FCS.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operand is optional:

"ALL"

Enables the MS topology discovery service throughout the fabric. This operand must be enclosed in double quotation marks.

Examples To enable the MS topology discovery service on the local switch:

```
switch:admin> mstdenable

Request to enable MS Topology Discovery Service
in progress...
done.
*MS Topology Discovery enabled locally.
```

To enable MS topology discovery on all switches in the fabric:

```
switch:admin> mstdenable "ALL"

Request to enable MS Topology Discovery Service
in progress...
done.
*MS Topology Discovery enabled locally.
*MS Topology Discovery Enable Operation Complete!!
```

SEE ALSO [msTdDisable](#), [msTdReadConfig](#)

msTdReadConfig

Displays the status of The Management Server (MS) topology discovery service.

Synopsis `mstdreadconfig`

Description Use this command to check whether or not the management server topology discovery service is enabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the status of the topology discovery service:

```
switch:admin> mstdreadconfig

*MS Topology Discovery is enabled.
```

SEE ALSO [msCapabilityShow](#), [msConfigure](#), [msPIMgmtActivate](#), [msPIMgmtDeactivate](#), [msTdDisable](#), [msTdEnable](#)

myId

Displays the current login session details.

Synopsis `myid`

Description Use this command to display the status of the system and the login session details. This includes IPv4 or IPv6 addresses associated with the login session.

The login session gives details of the following:

- CP/switch (or console/serial port) used to log in.
- The IP address of the current login session for Telnet or the name of the current console port or the serial port (if modem login used).
- The current CP mode (Active, Standby, or N/A).
- The current system status (Redundant, Nonredundant, or N/A).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display current login information:

```
switch:admin> myid
Current Switch: switch
Session Detail: switch (123.123.123.123) Active Redundant
```

```
switch:admin> myid
Current Switch: sw0
Session Detail: sw0 (123.123.123.123) N/A HA-Status-N/A
```

SEE ALSO [version](#)

nbrStateShow

Displays the state of FSPF neighbors.

Synopsis `nbrstateshow [slot][port]`

Description Use this command to display information about fabric shortest path first (FSPF) neighbors to the local switch or information about a neighbor to a specified port. FSPF defines a neighbor as a remote E_Port interface that is directly attached to the local switch. However, if ports are trunked, the command displays data only about the trunk master.

This command displays the following fields:

| | |
|------------------------|--|
| Local Domain ID | Domain ID of the local switch. |
| Local Port | E_Port interface on the local switch. This value is typically equal to the Index field reported in the switchShow command. |
| Domain | Domain ID of the remote switch. |
| Remote Port | E_Port interface on the remote switch. |
| State | State of the neighbor. The neighbor can be in one of the following five states: |
| 0 | NB_ST_DOWN - The neighbor is down. |
| 1 | NB_ST_INIT - The neighbor is initializing. |
| 2 | NB_ST_DB_EX - The neighbor and the switch are exchanging data from their Link State Records (LSR) databases. |
| 3 | NB_ST_DB_ACK_WT -The neighbor is waiting for the switch to acknowledge the LSR database. |
| 4 | NB_ST_DB_WT - The LSR Database is in waiting state; synchronization is in process. |
| 5 | NB_ST_FULL - The neighbor is in the last, finishing state. The E_Port can route frames only if the neighbor is in full state. |
| Sub-State | Substate of the port. The port can be in one of the following 12 substates: |
| INIT | The port is initializing. |

| | |
|-----------------------|--|
| ROUTABLE | The port is ready to receive frames. |
| ROUTABLE_WAIT | The port is waiting for notification from the neighbor. |
| ROUTABLE_SEND | The port is preparing for local route update; negotiates for the necessary locks before updating the routes. |
| ROUTABLE_TIMER | The request for one or more locks failed; delay for a short interval before retrying. |
| DONE | The port is online and in use. |
| DECOM_START | The decommissioning request is sent to the neighbor and waiting for a response. |
| DECOM_WAIT | Waiting for neighbor to remove routes using the interswitch link (ISL). |
| DECOM_SEND | Preparing to update routes of the local port; negotiates for the necessary before updating the routes. |
| DECOM_TIMER | The request for one or more locks failed; delay for a short interval before retrying. |
| DECOM_DONE | The local routes are updated; waiting for acknowledgement from the neighbor. |
| DECOM_BLOCK | The local routes are updated; the neighbor sent acknowledgement. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the port to display, followed by a slash (/).

port

Specify the port number to display, relative to its slot for bladed systems. Use **switchShow** to list valid ports. This operand is optional; if omitted, all neighbor states are displayed.

Examples To display information about a neighbor directly connected to the local switch:

```
switch:user> nbrstateshow 2/0
Local Domain ID: 1

Local Port   Domain   Remote Port   State   Sub-State
-----
          16         2           48     NB_ST_FULL  DONE
```

SEE ALSO [interfaceShow](#)

nbrStatsClear

Resets FSPF interface counters.

Synopsis `nbrstatsclear [slot][port]`

Description Use this command to reset the counters of fabric shortest path first (FSPF) frames transmitted and received on all interswitch links (ISLs) or on a specified ISL. Use this command without operands to reset counters on all interfaces. Use **interfaceShow** to view the FSPF counters.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the port to display, followed by a slash (/).

port

Specify the port number to display, relative to its slot for bladed systems. Use **switchShow** to list valid ports. This operand is optional; if omitted, FSPF statistics are reset.

Examples To display the counters on a port:

```
switch:admin> interfaceshow 1/0

idbP          = 0x10050a38

Interface 0 data structure:

nghbP        = 0x1004ce68
ifNo         = 0
masterPort   = 0 (self)
defaultCost  = 500
cost         = 500
delay        = 1

(output truncated)

nCmdAcc      = 37
nInvCmd      = 0
nHloIn       = 10
nInvHlo      = 0
nLsuIn       = 17
nLsaIn       = 10
attHloOut    = 11
nHloOut      = 11
attLsuOut    = 12
nLsuOut      = 12
attLsaOut    = 17
nLsaOut      = 17
```

To reset the counters on a port:

```
switch:admin> nbrstatsclear 1/0
```

To verify the changes:

```
switch:admin> interfaceshow 1/0

idbP          = 0x10050a38

Interface 0 data structure:

nghbP        = 0x1004ce68
ifNo         = 0
masterPort   = 0 (self)
defaultCost  = 500
cost         = 500

(output truncated)
```

SEE ALSO [interfaceShow](#), [portShow](#), [switchShow](#)

nodeFind

Displays all device Name Server (NS) entries matching a given WWN, device PID, or alias.

Synopsis `nodefind WWN | PID | ALIAS`

Description Use this command to display the NS information for all devices in the fabric that have either a port world wide name (WWN) or a node WWN matching the given WWN; or have a device PID matching the given PID; or have a defined configuration alias to which the device belongs matching the given alias.

If there is no device matching the given WWN, PID, or alias, the message *"No device found"* is displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

WWN | PID | ALIAS

Specify the WWN, device PID, or alias that can be used to match the real device's data. The WWN must be specified as a hexadecimal value or a colon-separated value with eight fields, each consisting of one or two hexadecimal numbers between 0 and ff. The PID usually begins with 0x or 0X. If the PID does not begin with 0x or 0X, it is interpreted as an alias. If an alias is not found, the argument is checked for other PIDs.

Examples To display all the device information matching the WWN "20:00:00:e0:8b:01:ce:d3":

```
switch:user> nodefind 200000e08b01ced3
Remote:
  Type Pid      COS      PortName      NodeName
  NL 020eef;3;20:00:00:e0:8b:01:ce:d3;20:00:00:e0:8b:01:ce:d3;
      Fabric Port Name: 20:0e:00:60:69:51:0b:ba
      Permanent Port Name: 20:00:00:e0:8b:01:ce:d3
Device type: Physical Target
Port Index: 14
Share Area: No
Device Shared in Other AD: No
Aliases:
```

To display all the device information matching the PID "0x020eef":

```
switch:user> nodefind 0x020eef
Remote:
  Type Pid      COS      PortName      NodeName
  NL 020eef; 3;20:00:00:e0:8b:01:ce:d3;20:00:00:e0:8b:01:ce:d3;
      Fabric Port Name: 20:0e:00:60:69:51:0b:ba
      Permanent Port Name: 20:00:00:e0:8b:01:ce:d3
Device type: Physical Target
Port Index: 14
Share Area: No
Device Shared in Other AD: No
Aliases:

No match:
```

To display device information for a string for which there is no match:

```
switch:user> nodefind abcd
No device found.
```

To display all the device information matching the alias "a320":

```
switch:user> nodefind a320
Local:
  Type Pid      COS      PortName      NodeName      SCR
NL 0314d9;3;22:00:00:04:cf:5d:dc:2d;20:00:00:04:cf:5d:dc:2d; 0
   FC4s: FCP [SEAGATE ST318452FC      0001]
   Fabric Port Name: 20:14:00:60:69:80:04:79
   Permanent Port Name: 22:00:00:04:cf:5d:dc:2d
   Device type: Physical Target
   Port Index: 20
   Share Area: No
   Device Shared in Other AD: No
   Aliases: a320
NL 0314d6; 3;22:00:00:04:cf:9f:78:7b;20:00:00:04:cf:9f:78:7b; 0
   FC4s: FCP [SEAGATE ST336605FC      0003]
   Fabric Port Name: 20:14:00:60:69:80:04:79
   Permanent Port Name: 22:00:00:04:cf:9f:78:7b
   Device type: Physical Target
   Port Index: 20
   Share Area: No
   Device Shared in Other AD: No
   Aliases: a320
NL 0314d5;3;22:00:00:04:cf:9f:7d:e0;20:00:00:04:cf:9f:7d:e0; 0
   FC4s: FCP [SEAGATE ST336605FC      0003]
   Fabric Port Name: 20:14:00:60:69:80:04:79
   Permanent Port Name: 22:00:00:04:cf:9f:7d:e0
   Device type: Physical Target
   Port Index: 20
   Share Area: No
   Device Shared in Other AD: No
   Aliases: a320
NL 0314d4;3;22:00:00:04:cf:9f:26:7e;20:00:00:04:cf:9f:26:7e; 0
   FC4s: FCP [SEAGATE ST336605FC      0003]
   Fabric Port Name: 20:14:00:60:69:80:04:79
   Permanent Port Name: 22:00:00:04:cf:9f:26:7e
   Device type: Physical Target
   Port Index: 20
   Share Area: No
   Device Shared in Other AD: No
   Aliases: a320
```

SEE ALSO [aliShow](#), [nsAllShow](#), [nsCamShow](#), [nsShow](#)

nsAliasShow

Displays local Name Server (NS) information, with aliases.

Synopsis `nsaliasshow [-r][-t] | -domain domain_id | --help`

Description Use this command to display local name server information with the added feature of displaying the defined configuration aliases to which the device belongs. If there are no defined configuration aliases for that device, no alias is displayed. If there is no information in this switch, the following message is displayed: "There is no entry in the Local Name Server." Use the **-domain *domainid*** option to display the remote device details for a specific domain in the fabric.

The information displayed for each device is the same that is displayed by the **nsShow** command with the exception of the additional display of the alias to which the device belongs. Refer to the **nsShow** help page for a description of these displays. Use **nsCamShow** to display information from all switches.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

| | |
|--------------|--|
| -r | Replaces the time-to-live (TTL) attribute output with state change registration (SCR) information. This value indicates what type of RSCN a device registers to receive. Values include the following: |
| SCR=0 | Reserved. |
| SCR=1 | Fabric detected registration. Register to receive all RSCN requests issued by the fabric controller for events detected by the fabric. |
| SCR=2 | NX_Port detected registration. Register to receive all RSCN requests issued for events detected by the affected NX_Port. |
| SCR=3 | Register to receive all RSCN requests issued. The RSCN request returns all effected N_Port_ID pages. |

-t

Displays the device type. The device type is defined in terms of two attributes. The first attribute indicates the origination of the device as one of the following:

- Physical -The device is connected to the NX_Port, using FLOGI to log in to the switch.
- Virtual -The device is contrived by the switch.
- NPV -The device is connected to the NX_Port, using FDISC to log in to the switch.
- iSCSI Device is connected to the iSCSI port.

The second attribute indicates the role of the device. Valid role attributes include the following:

- Unknown (initiator/target) - Device role is not detected
- Initiator - An iSCSI initiator.
- Target - An iSCSI target.
- Initiator+Target - Both an iSCSI initiator and an iSCSI target.

-domain *domain_id*

Displays the remote device details for a specific domain.

--help

Displays the command usage.

Examples

To display local NS information with aliases:

```
switch:user> nsaliasshow
{
Type Pid      COS   PortName      NodeName      TTL(sec)
N 010600;3;20:06:00:05:1e:38:81:71;10:00:00:05:1e:38:81:71; na
Fabric Port Name: 20:06:00:05:1e:7a:7a:00
Permanent Port Name: 20:06:00:05:1e:38:81:71
Port Index: 6
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Aliases: MyAlias1 MyAlias2
N 010601;3;23:0d:00:05:1e:38:81:71;50:00:51:e3:88:17:10:0d; na
FC4s: FCP
PortSymb: [44] "Brocade Ioblaster Port Entity #00,pid#10601."
Fabric Port Name: 20:06:00:05:1e:7a:7a:00
Permanent Port Name: 20:06:00:05:1e:38:81:71
Port Index: 6
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Aliases:
N 010602;3;10:00:00:00:00:00:00:01;10:00:00:00:00:00:00:01; na
FC4s: FCP
PortSymb: [41] "Brocade Ioblaster Initiator#00,pid#10602."
Fabric Port Name: 20:06:00:05:1e:7a:7a:00
Permanent Port Name: 20:06:00:05:1e:38:81:71
Port Index: 6
Share Area: No
```

```

    Device Shared in Other AD: No
    Redirect: No
    Partial: No
    Aliases: DeviceAlias
  The Local Name Server has 3 entries }

```

To display local NS information with aliases with the **-r** option:

```

switch:user> nsaliasshow -r
{
  Type Pid  COS      PortName          NodeName          SCR
N 010600;3;20:06:00:05:1e:38:81:71;10:00:00:05:1e:38:81:71; 0
  Fabric Port Name: 20:06:00:05:1e:7a:7a:00
  Permanent Port Name: 20:06:00:05:1e:38:81:71
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases: MyAlias1 MyAlias2
N 010601; 3;23:0d:00:05:1e:38:81:71;50:00:51:e3:88:17:10:0d; 3
  FC4s: FCP
  PortSymb: [44] "Brocade Ioblaster Port Entity #00,pid#10601."
  Fabric Port Name: 20:06:00:05:1e:7a:7a:00
  Permanent Port Name: 20:06:00:05:1e:38:81:71
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases:
N 010602;3;10:00:00:00:00:00:00:01;10:00:00:00:00:00:00:01; 3
  FC4s: FCP
  PortSymb: [41] "Brocade Ioblaster Initiator#00,pid#10602."
  Fabric Port Name: 20:06:00:05:1e:7a:7a:00
  Permanent Port Name: 20:06:00:05:1e:38:81:71
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases: DeviceAlias
  The Local Name Server has 3 entries }

```

To display local NS information with aliases with the **-r** and **-t** options:

```

switch:admin> nsaliasshow -r -t
{
  Type Pid  COS      PortName          NodeName          SCR
N 010600; 3;20:06:00:05:1e:38:81:71;10:00:00:05:1e:38:81:71; 0
  Fabric Port Name: 20:06:00:05:1e:7a:7a:00
  Permanent Port Name: 20:06:00:05:1e:38:81:71
  Device type: Physical Unknown(initiator/target)
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases: MyAlias1 MyAlias2

```

```

N 010601;3;23:0d:00:05:1e:38:81:71;50:00:51:e3:88:17:10:0d; 3
  FC4s: FCP
  PortSymb: [44] "Brocade Ioblaster Port Entity #00,pid#10601."
  Fabric Port Name: 20:06:00:05:1e:7a:7a:00
  Permanent Port Name: 20:06:00:05:1e:38:81:71
  Device type: NPIV Unknown(initiator/target)
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases:
N 010602; 3;10:00:00:00:00:00:00:01;10:00:00:00:00:00:00:01; 3
  FC4s: FCP
  PortSymb: [41] "Brocade Ioblaster Initiator#00,pid#10602."
  Fabric Port Name: 20:06:00:05:1e:7a:7a:00
  Permanent Port Name: 20:06:00:05:1e:38:81:71
  Device type: NPIV Initiator
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases: DeviceAlias
The Local Name Server has 3entries }

```

To display remote device details for a specific domain:

```

switch:admin> nsaliasshow -domain 92
{
  Type Pid    COS    PortName                               NodeName
N 5c1000; 3;23:06:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:06;
  FC4s: FCP
  PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#03."
  Fabric Port Name: 20:10:00:05:1e:53:e3:8a
  Permanent Port Name: 23:06:00:05:1e:53:e3:8a
  Port Index: 16
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases: MyAlias1
N 5c1200;3;23:04:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:04;
  FC4s: FCP
  PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#02."
  Fabric Port Name: 20:12:00:05:1e:53:e3:8a
  Permanent Port Name: 23:04:00:05:1e:53:e3:8a
  Port Index: 18
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  Aliases: MyAlias1
N 5c1300; 3;23:02:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:02;
  FC4s: FCP
  PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#01."
  Fabric Port Name: 20:13:00:05:1e:53:e3:8a
  Permanent Port Name: 23:02:00:05:1e:53:e3:8a
  Port Index: 19
  Share Area: No

```

nsAliasShow

```
Device Shared in Other AD: No
Redirect: No
Partial: No
Aliases:
N 5c1700; 3;23:00:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:00;
FC4s: FCP
PortSymb: [34] "Brocade DPC Entity-Slot#00,DPC#00."
Fabric Port Name: 20:17:00:05:1e:53:e3:8a
Permanent Port Name: 23:00:00:05:1e:53:e3:8a
Port Index: 23
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Aliases:
The Remote Name Server has 4 entries }
```

SEE ALSO [nsAllShow](#), [nsShow](#), [switchShow](#)

nsAllShow

Displays global name server information.

Synopsis `nsallshow [type]`

Description Use this command to display the 24-bit Fibre Channel addresses of all devices in all switches in the fabric. When used with the type operand, the command displays only devices of the specified FC-4 type. FC-4 type codes are referenced in the *Fibre Channel Framing and Signaling* (FC-FS) standards documentation (see "TYPE codes - FC-4"). When used without operand, all devices are displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

type

Specifies the FC-4 type code to filter the command output. Valid values are 0 to 255. Only the following two FC-4 device type codes are applicable to the Brocade environment:

- 8 - FCP type device
- 5 - FC-IP type device

For all other codes, entries are summarized in the format "*x* ports supporting FC4 *type*", where *x* is the number of ports and *type* is the user-specified FC-4 type code in hexadecimal format.

Examples To display all devices in the fabric, followed by all type 8 (SCSI-FCP) devices, and all type 5 (SCSI-FCIP) devices:

```
switch:admin> nsAllShow
{
  011000 011200 0118e2 0118e4 0118e8 0118ef 021200 021300
  0214e2 0214e4 0214e8 0214ef
  12 Nx_Ports in the Fabric }
```

```
switch:admin> nsAllShow 8
{
  0118e2 0118e4 0118e8 0118ef 0214e2 0214e4 0214e8 0214ef
  8 FCP Ports }
```

```
switch:admin> nsAllShow 5
{
  011200 021200
  2 FC-IP Ports }
```

To display a device type of 255:

```
switch:admin> nsAllShow 255
{
  010100 020a00
  2 Ports supporting FC4 0xff }
```

SEE ALSO [nsShow](#), [switchShow](#)

nsCamShow

Displays information about remote devices in the Name Server (NS) cache.

Synopsis `nscamshow [-t]`

Description Use this command to display the local NS cache information about the devices discovered in the fabric by the NS cache manager.

If the NS cache manager does not discover new switches or new devices in the fabric, the command displays the message "No Entry is found!"

For each discovered remote switch, this command displays the following information:

Switch entry for *N*

Displays the remote domain ID for the switch.

state

Displays one of the following values:

known

The local domain is aware of all the devices from this remote domain.

unknown

The local domain is unaware of devices from this remote domain.

ERROR

The information for this remote domain is unreliable.

rev

Fabric OS firmware version of the remote switch. For switches running firmware other than Fabric OS, a string of question marks ("????") is displayed.

owner

Displays the owner of the NSCAM database entry. The value displayed is a domain ID (domain address). For example, 0xffc02 indicates domain 2 and is the domain on which the command has been executed. This is the local domain and the information is stored locally by this switch.

cap_available

Each switch in the fabric exchanges information regarding its capabilities (for example, firmware level, feature support, etc.). When the **cap_available** value is 1, it indicates that the local domain has received the capabilities of the remote domain that is being displayed.

When the value is 0 capability information has not been received.

The remaining information displayed for each device is the same that is displayed by the **nsShow** command. Refer to the **nsShow** help page for a description of these displays.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operand is optional:

| | |
|-----------------------------------|---|
| -t | Displays the device type. The device type is defined in terms of two attributes. The first attribute indicates the origination of the device as one of the following: |
| Physical | The device is connected to the NX_Port, using FLOGI to log in to the switch. |
| Virtual | The device is contrived by the switch. |
| NPV | The device is connected to the NX_Port, using FDISC to log in to the switch. |
| iSCSI | Device is connected to the iSCSI port. |
| iSCSI | The second attribute indicates the role of the device. Valid role attributes include the following: |
| Unknown (initiator/target) | Device role is not detected |
| Initiator | An iSCSI initiator. |
| Target | An iSCSI target. |
| Initiator+Target | Both an iSCSI initiator and an iSCSI target. |

Examples To display all switch and device entries discovered by the Name Server in the fabric:

```
switch:user> nscamshow
nscam show for remote switches:
Switch entry for 92
state rev owner cap_available
known v640 0xffffc23 1
Device list: count 8
Type Pid COS PortName NodeName
N 5c1000; 3;23:06:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:06;
FC4s: FCP
PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#03."
Fabric Port Name: 20:10:00:05:1e:53:e3:8a
Permanent Port Name: 23:06:00:05:1e:53:e3:8a
Port Index: 16
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
N 5c1200;3;23:04:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:04;
FC4s: FCP
```

```

PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#02."
Fabric Port Name: 20:12:00:05:1e:53:e3:8a
Permanent Port Name: 23:04:00:05:1e:53:e3:8a
Port Index: 18
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Port Properties: SIM Port
N 5c1300; 3;23:02:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:02;
FC4s: FCP
PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#01."
Fabric Port Name: 20:13:00:05:1e:53:e3:8a
Permanent Port Name: 23:02:00:05:1e:53:e3:8a
Port Index: 19
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Port Properties: SIM Port
N 5c1700; 3;23:00:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:00;
FC4s: FCP
PortSymb: [34] "Brocade DPC Entity-Slot#00,DPC#00."
Fabric Port Name: 20:17:00:05:1e:53:e3:8a
Permanent Port Name: 23:00:00:05:1e:53:e3:8a
Port Index: 23
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
Port Properties: SIM Port
(output truncated)

```

To display the output with the `-t` option:

```

switch:user> nscamshow -t
Switch entry for 92
state rev owner cap_available
known v640 0xffffc23 1
Device list: count 8
Type Pid COS PortName NodeName
N 5c1000;3;23:06:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:06;
FC4s: FCP
PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#03."
Fabric Port Name: 20:10:00:05:1e:53:e3:8a
Permanent Port Name: 23:06:00:05:1e:53:e3:8a
Device type: Virtual Unknown(initiator/target)
Port Index: 16
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
N 5c1200; 3;23:04:00:05:1e:53:e3:8a;50:00:51:e5:3e:38:a0:04;
FC4s: FCP
PortSymb: [45] "Brocade VDPC Entity-Slot#00,DPC#00,Entity#02."
Fabric Port Name: 20:12:00:05:1e:53:e3:8a
Permanent Port Name: 23:04:00:05:1e:53:e3:8a
Device type: Virtual Unknown(initiator/target)
Port Index: 18
Share Area: No

```



```
Device Shared in Other AD: No
Redirect: No
Partial: No
Port Properties: SIM Port
(output truncated)
```

SEE ALSO [nsAllShow](#), [nsAliasShow](#), [nsShow](#), [switchShow](#)

nsDevLog

Manages device history logging.

Synopsis `nsdevlog --show` *[[-slot slot] -port port |
-pid pid | -wwpn *wwpn* | -wwnn *nwwn* | -event *event*]*

`nsdevlog --enable`

`nsdevlog --disable`

`nsdevlog --clear`

`nsdevlog --help`

Description Use this command to manage Name Server (NS) device logging. By default, logging is enabled. The NS records the PortIndex, PID, world wide node name (WWNN), world wide port name (WWPN) and the event (login and logout).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

`--show options`

Displays the device logs.

`-slot slot -port port`

Displays device logs for the specified port or slot/port.

`-pid pid`

Displays device logs for the specified PID.

`-wwpn ppwwn`

Displays device logs for the specified port WWN.

`-wwnn nwwn`

Displays device logs for the specified node WWN.

| | |
|---------------------|---|
| -event event | Displays device logs for the specified event. The events are the following: <ul style="list-style-type: none"> • "Register": Entry created by explicit NS registration. • "Deregister": Entry deleted by explicit NS deregistration. • "Device Add": Entry created based on UPD Area. • "Device Del": Entry deleted based on UPD Delete. • "Device Add (Query)": Entry created by NS query from the device. • "Port Del": Entries deleted based on DEL Area. • "Dup WWN": Entry deleted based on Duplicate WWN. • "Switch Offline": Entries deleted due to switch offline. • "FPORT Entry": Entry created by F-Port SCN. |
| --enable | Enables NS device history logging. |
| --disable | Disables NS device history logging. |
| --clear | Clears the NS device history logs. |
| --help | Displays the command usage. |

Examples

To enable device logging:

```
switch:admin> nsdevlog --enable
Device History Logging enabled
```

To disable device logging:

```
switch:admin> nsdevlog --disable
Device History Logging disabled
```

To display the NS device log for an event:

```
switch:admin> nsdevlog --show -event "FPORT Entry"
date/time          slot/port PID      Port WWN          \
      Node WWN          event
=====
Wed Jun 27 11:24:33.608  2/12  0x034c00  30:11:00:05:1e:0f:9c:6a \
      20:11:00:05:1e:0f:9c:6a  FPORT Entry
Wed Jun 27 11:24:33.881  2/11  0x034b00  30:1c:00:05:1e:0f:9c:6a \
      20:1c:00:05:1e:0f:9c:6a  FPORT Entry
Wed Jun 27 11:24:34.028  7/28  0x039c00  20:0c:00:05:1e:06:c3:1a \
      10:00:00:05:1e:06:c3:1a  FPORT Entry
Wed Jun 27 11:24:34.048  7/29  0x039d00  20:08:00:05:1e:06:c3:1a \
      10:00:00:05:1e:06:c3:1a  FPORT Entry
Total number of Logged entries = 133
Total number of Entries displayed = 4
Max number of entries: 65536
Device Logging is currently disabled
```

nsDevLog

To clear device logs:

```
switch:admin> nsdevlog --clear  
Device History logs cleared
```

To display device logs:

```
switch:admin> nsdevlog --show  
Total number of Logged entries = 0  
Total number of Entries displayed = 0  
Max number of entries: 65536  
Device Logging is currently disabled
```

SEE ALSO [nsAliasShow](#), [nsAllShow](#), [nsCamShow](#), [nsShow](#), [nsZoneMember](#)

nsShow

Displays local Name Server (NS) information.

Synopsis `nsshow [-r][-t]`

Description Use this command to display local NS information about devices connected to a switch. If no information is available for the switch, the command displays the following message: "There is no entry in the Local Name Server."

Use **nsAllShow** to display NS information for all switches in the fabric.

Each line of output displays the following information:

| | |
|----------------------------|---|
| Type | Displays U for unknown, N for N_Port, NL for NL_Port. |
| PID | Displays the 24-bit Fibre Channel address of the device. |
| COS | Displays the Class of Service levels supported by the device. This can be class 1, class 2, or class 3. The command displays only the numeric values, 1, 2, and/or 3. A device can support multiple CoS levels. |
| PortName | Displays the device port world wide name (WWN). |
| NodeName | Displays the device node WWN. |
| TTL | Displays the time-to-live, in seconds, for cached entries or NA (not applicable) if the entry is local. |
| SCR | Displays the state change registration of the device. This field displays only with the -r option. |
| Device type | Displays the device type if -t is specified. |
| Fabric Port Name | Displays the F_Port WWN to which the N_Port connects. |
| Permanent Port Name | Displays the physical N_Port or NL_Port WWN. |
| Port Index | Displays the index number of the physical N_Port to which the device connects. |
| Share Area | Displays "Yes" if the port shares an area; otherwise displays "No". |

Redirect

Displays "Yes" if the device is involved in frame redirection; otherwise displays "No". The device involved in frame redirection is specified as either "virtual" , "host", or "target".

Partial

Displays "Yes" if the device entry is incomplete; otherwise displays "No". Devices that are incomplete are displayed by the **nsShow** and **nsCamShow** commands, and have routing established, but are not considered during device discovery (for example, during FC-GS Name Server Queries).

LSAN

Displays "Yes" if the device is currently part of an active LSAN zone; otherwise displays "No".

Device link speed

Displays the link speed of the device, for example: 4G, 8G, 16G and 32G. Link speed of end devices that log in to the edge fabric through F_Port trunk display the bandwidth of the F_Port trunk. For example, if an 8 Gbps host attached to AG login to fabric through a 32 Gbps F_Port trunk, the device link speed for the 8 Gbps host is displayed as 32G.

Port Properties

Displays a list of port properties, for example: "SIM Port" for SIM ports. If no properties are assigned, the field is not displayed.

The following information is displayed only if the device has registered the information (for example, the switch automatically registers SCSI inquiry data for FCP target devices):

- FC4s supported
- IP address
- Port and node symbolic names
- Fabric Port Name. This is the WWN of the port on the switch to which the device is physically connected.
- Port address or port IP address

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following optional operands:

-r

Replaces the time-to-live (TTL) attribute output with state change registration (SCR) information. This value indicates what type of RSCN a device registers to receive. Values include the following:

SCR=0x00000000

Reserved.

SCR=0x00000001

Fabric detected registration. Register to receive all RSCN requests issued by the fabric controller for events detected by the fabric.

SCR=0x00000002

NX_Port detected registration. Register to receive all RSCN requests issued for events detected by the affected NX_Port.

SCR=0x00000003

Register to receive all RSCN requests issued. The RSCN request returns all effected N_Port_ID pages.

The following SCR values use the 4th byte of the SCR value and can coexist with the SCR values of 0x00000000, 0x00000001, 0x00000002, and 0x00000003.

SCR=0x01000000

Principal switch RSCN registration. Register to receive proprietary RSCN (principal switch change).

SCR=0x02000000

AG_QOS RSCN registration.

-t

Displays the device type. The device type is defined in terms of two attributes. The first attribute indicates the origination of the device as one of the following:

Physical

The device is connected to the NX_Port, using FLOGI to log in to the switch.

Virtual

The device is contrived by the switch.

NPV

The device is connected to the NX_Port, using FDISC to log in to the switch.

iSCSI

The device is connected to the iSCSI port.

The second attribute indicates the role of the device. Valid role attributes include the following:

Unknown (initiator/target)

Device role is not detected

Initiator

An iSCSI initiator.

Target

An iSCSI target.

Initiator+Target

Both an iSCSI initiator and an iSCSI target.

Examples

To display local NS information:

```
switch:user> nsshow
{
  Type Pid COS      PortName          NodeName          TTL(sec)
  N 010100;3;21:00:00:e0:8b:13:08:10;20:00:00:e0:8b:13:08:10;na
    FC4s: FCP
    NodeSymb: [41] "QLA2340 FW:v3.03.06 DVR:v9.0.0.2 (w32 IP)"
    Fabric Port Name: 20:01:00:05:1e:34:00:70
    Permanent Port Name: 21:00:00:e0:8b:13:08:10
    Port Index: 1
    Share Area: No
    Device Shared in Other AD: No
}
```

```

Redirect: No
Partial: No
LSAN: No
Device link speed: 8G
N 010e00;3;21:00:00:e0:8b:12:8a:be;20:00:00:e0:8b:12:8a:be;na
FC4s: FCP
NodeSymb: [41] "QLA2340 FW:v3.03.06 DVR:v9.0.0.2 (w32 IP)"
Fabric Port Name: 20:0e:00:05:1e:34:00:70
Permanent Port Name: 21:00:00:e0:8b:12:8a:be
Port Index: 14
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
LSAN: No
Device link speed: 4G
Port Properties: SIM Port
The Local Name Server has 2 entries }

```

To display local name server information with the `-r` option.

```

switch:user> nsshow -r
Type Pid COS PortName NodeName SCR
N 010100;3;21:00:00:e0:8b:13:08:10;20:00:00:e0:8b:13:08:10; \
0x01000003
FC4s: FCP
NodeSymb: [41] "QLA2340 FW:v3.03.06 DVR:v9.0.0.2 (w32 IP)"
Fabric Port Name: 20:01:00:05:1e:34:00:70
Permanent Port Name: 21:00:00:e0:8b:13:08:10
Port Index: 1
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
LSAN: No
Device link speed: 4G
N 010e00;3;21:00:00:e0:8b:12:8a:be;20:00:00:e0:8b:12:8a:be; \
0x01000001
FC4s: FCP
NodeSymb: [41] "QLA2340 FW:v3.03.06 DVR:v9.0.0.2 (w32 IP)"
Fabric Port Name: 20:0e:00:05:1e:34:00:70
Permanent Port Name: 21:00:00:e0:8b:12:8a:be
Port Index: 14
Share Area: No
Device Shared in Other AD: No
Redirect: No
Partial: No
LSAN: No
Device link speed: 4G
Port Properties: SIM Port
The Local Name Server has 2 entries }

```


To display local name server information with `-r` and `-t` options.

```
switch:user> nsshow -r -t
{
  Type Pid COS      PortName          NodeName          SCR
N 010100; 3;21:00:00:e0:8b:13:08:10;20:00:00:e0:8b:13:08:10; \
  0x00000001
  FC4s: FCP
  NodeSymb: [41] "QLA2340 FW:v3.03.06 DVR:v9.0.0.2 (w32 IP)"
  Fabric Port Name: 20:01:00:05:1e:34:00:70
  Permanent Port Name: 21:00:00:e0:8b:13:08:10
  Device type: Physical Initiator
  Port Index: 1
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  LSAN: No
  Device link speed: 4G
  Port Properties: SIM Port
N 010e00;3;21:00:00:e0:8b:12:8a:be;20:00:00:e0:8b:12:8a:be; \
  0x00000001
  FC4s: FCP
  NodeSymb: [41] "QLA2340 FW:v3.03.06 DVR:v9.0.0.2 (w32 IP)"
  Fabric Port Name: 20:0e:00:05:1e:34:00:70
  Permanent Port Name: 21:00:00:e0:8b:12:8a:be
  Device type: Physical Initiator
  Port Index: 14
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  Partial: No
  LSAN: No
  Device link speed: 4G
  Port Properties: SIM Port
The Local Name Server has 2 entries }
```

SEE ALSO [nsAllShow](#), [nsAliasShow](#), [nsCamShow](#), [switchShow](#)

nsZoneMember

Displays the information on online devices zoned with a specified device.

Synopsis `nszonemember pid | wwn`

`nszonemember -a | -n | -u`

`nszonemember [-domain domain] [-index index]`

`nszonemember --help`

Description Use this command to display information on all online devices zoned with the specified device. The device can be specified by WWN or Port ID (PID). Use this command with the `-u` option to display all unzoned devices in the entire fabric. Use the `-a` option to display online zoned device data for each local device. Use the `-domain` and `-index` options to display zoned device data for a device (either local or remote) in the fabric with the specified domain and port index combination. If a domain is not specified, device data for a local device with the specified port index is displayed.

The command output displays the following information:

| | |
|----------------------------------|--|
| Type | U - known, N - N_Port, NL - NL_Port. |
| Pid | The 24-bit Fibre Channel address. |
| COS | A list of classes of service supported by the device. |
| PortName | The device's port world wide name (WWN). |
| NodeName | The device's node WWN. |
| Permanent Port Name | The physical N_Port or NL_Port WWN. |
| DeviceType | The device type. |
| Port Index | The index of the port to which the device is attached. |
| Shared Area | Whether or not the device shares an area with other devices. |
| Device Shared in Other AD | Whether or not the device is shared in other Admin Domains. |

Additional lines may display if the device has registered any of the following information (the switch automatically registers SCSI inquiry data for FCP target devices):

- FC4 supported
- IP address (node)
- IPA
- port and node symbolic name (local device only)
- fabric port name
- hard address or port IP address

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------------------|--|
| <i>PID</i> <i>WWN</i> | Specifies the port ID or WWN of the device for which to display zoned devices. |
| -a | Displays each local device's online zoned device data, including the device PID and zone alias. |
| -n | Displays each local device's online zoned device data including the device PID. |
| -u | Displays all unzoned devices in the entire fabric. The device data includes the device PID and zone alias. |
| -domain <i>domain</i> | Displays zoned device data for a device (either local or remote) with the specified domain. |
| -index <i>index</i> | Displays zoned device data for a device (either local or remote) with the specified port index. |
| --help | Displays the command usage. |

Examples To display information about all the online devices zoned with the given device:

```
switch:admin> nszonemember Ox0416e2
3 local zoned members:

Type Pid COS   PortName                               NodeName                               SCR
NL 041901;2,3;10:00:00:00:c9:26:0e:ae;20:00:00:00:c9:26:0e:ae; 3
Fabric Port Name: 20:09:00:60:69:50:06:78
Permanent Port Name: 10:00:00:00:c9:26:0e:ae
Device type: Physical Initiator
NL 0416e2;3;22:00:00:20:37:d9:6b:b3;20:00:00:20:37:d9:6b:b3; 0
FC4s: FCP [SEAGATE ST318304FC 0005]
Fabric Port Name: 20:06:00:60:69:50:06:78
Permanent Port Name: 22:00:00:20:37:d9:6b:b3
Device type: Physical Target
NL 0416e4; 3;22:00:00:20:37:d9:61:ac;20:00:00:20:37:d9:61:ac; 0
```

```

FC4s: FCP [SEAGATE ST318304FC      0005]
Fabric Port Name: 20:06:00:60:69:50:06:78
Permanent Port Name: 22:00:00:20:37:d9:61:ac
Device type: Physical Target

```

No remote zoned members

To display information about all the online devices zoned with the given WWN:

```

switch:admin> nszonemember 10:00:00:00:c8:23:0b:ad
3 local zoned members:

Type Pid  COS      PortName          NodeName          SCR
NL 041901;2,3;10:00:00:00:c9:26:0e:ae;20:00:00:00:c9:26:0e:ae;3
    Fabric Port Name: 20:09:00:60:69:50:06:78
    Permanent Port Name: 10:00:00:00:c9:26:0e:ae
    Device type: Physical Initiator
NL 0416e2; 3;22:00:00:20:37:d9:6b:b3;20:00:00:20:37:d9:6b:b3; 0
    FC4s: FCP [SEAGATE ST318304FC      0005]
    Fabric Port Name: 20:06:00:60:69:50:06:78
    Permanent Port Name: 22:00:00:20:37:d9:6b:b3
    Device type: Physical Target
NL 0416e4; 3;22:00:00:20:37:d9:61:ac;20:00:00:20:37:d9:61:ac; 0
    FC4s: FCP [SEAGATE ST318304FC      0005]
    Permanent Port Name: 22:00:00:20:37:d9:61:ac
    Device type: Physical Target

No remote zoned members

```

To display all the unzoned devices in the fabric:

```

switch:admin> nszonemember -u
Pid: 0xb01ea9;    Aliases: trimm32b_1
Pid: 0xb01eaa;    Aliases: trimm32b_2
Pid: 0xb01eab;    Aliases: trimm32b_3
Pid: 0xb01eac;    Aliases: trimm32b_4
Pid: 0xb01fad;    Aliases: trimm32a_5
Pid: 0xb01fae;    Aliases: trimm32a_6
Pid: 0xb01fb1;    Aliases: trimm32a_7
Pid: 0xb01fb2;    Aliases: trimm32a_8
Pid: 0xdc2800;    Aliases:
Totally 9 unzoned devices in the fabric.

```

To display each local device's online zoned device data:

```

switch:admin> nszonemember -a
Port: 4 Pid: 0xb00400    Aliases: ix360_131_201_6a
    Zoned Members: 2 devices
        Pid: 0xb00400    Aliases: ix360_131_201_6a
        Pid: 0xbalee8    Aliases: trimm101b_3

Port: 12 Pid: 0xb00c01    Aliases: dl360_130159a
    Zoned Members: 2 devices
        Pid: 0xb00c01    Aliases: dl360_130159a
        Pid: 0xbd1bef    Aliases: nstor4b_8

Port: 13 Pid: 0xb00d00    Aliases: ix360_131_196p5
    Zoned Members: 2 devices
        Pid: 0xb00d00    Aliases: ix360_131_196p5

```

```

    Pid: 0xe07d00    Aliases: hds9200_6p4 hds9200_6p4
Port: 14  Pid: 0xb00e00  Aliases: dl360_130251a dl360_130251a
    Zoned Members: 2 devices
    Pid: 0xb00e00    Aliases: dl360_130251a dl360_130251a
    Pid: 0xbalae4    Aliases: trimm100a_2

```

To display device data for a local device with the specified port index:

```

switch:admin> nszonemember -index 2
Port Index: 2  Pid: 0x015200
    Zoned Members: 2 devices
    Pid: 0x015200    Aliases: ali_b1
    Pid: 0x03a840    Aliases: ali_b1

```

To display zoned device data for a device (either local or remote) in the fabric with the specified domain and port index:

```

switch:admin> nszonemember -domain 3 -index 168
Port Index: 168  Pid: 0x03a840
    Zoned Members: 2 devices
    Pid: 0x03a840    Aliases: ali_z1
    Pid: 0x015200    Aliases: ali_z1

```

SEE ALSO [cfgShow](#), [nsCamShow](#), [nsShow](#)

nsZoneShow

Displays the zone names.

Synopsis

nszonestow -pid *pid*

nszonestow -wwn *wwn*

nszonestow -pid "*pid1,pid2*"

nszonestow -wwn "*wwn1,wwn2*"

Description

Use this command to display the zone names that a specified device or device-pair are part of. The devices can be zoned as part of regular zones, Traffic Isolation (TI) zones, or Frame Redirect (RD) zones.

- If you specify a single Port ID (PID) or world wide name (WWN), the command displays the names of the regular zones, RD zones, or TI zones that the devices belong to.
- If you specify a PID pair or WWN pair, the command displays the names of the common zones (regular, RD or TI zones) that the device pairs share.
- For devices zoned as part of an RD zone, the command displays the following information:
 - If you specify the PID or WWN pair for a given host and virtual target (VT) pair, the command displays the RD zone name.
 - If you specify the PID or WWN pair for a given virtual host (VI) and target pair, the command displays the RD zone name.
 - If you specify the PID or WWN for a given VI or VT, the command displays the RD zone name.
 - If you specify the PID or WWN pair for a given host and target pair, the command displays the regular zone name and the RD zone name.
 - If you specify the PID or WWN for a given host or target, the command displays the regular zone name and the RD zone name.

The host, target, VI and VT are specified while creating zones using the **zone** command. The RD zone names will start with the prefix 'red_O917'. You can use the **cfgshow** or the **zone --show** command to view the zone configuration information.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

-pid *pid*

Displays the zone names that the PID belongs to.

-wwn *wwn*

Displays the zone names that the WWN belongs to.

-pid "*pid1,pid2*"

Displays the zone names of the common zones where both PIDs are zoned together.

-wwn "*wwn1,wwn2*"

Displays the zone names of the common zones where both WWNs are zoned together.

Examples To display the zone names for a device with a specified PID:

```
switch:admin> nszonestow -pid 0x010200
Zone Names
=====
zn_test1
zn_test2
=====
```

To display the common zone for two devices specified by a WWN pair:

```
switch:admin> nszonestow \
wwn "20:08:00:05:1e:a3:01:d9,20:01:00:05:1e:a3:01:d9"
Zone Names
=====
zn_test2
=====
```

To display the RD zone name:

```
switch:admin> cfigshow
Defined configuration:
cfg: myHTcfg myHostTarget
cfg: r_e_d_i_r_c_fg
red_____base;
red_0917_00_3f_3f_3f_23_24_25_26_3f_3f_3f_30_32_00_00_00
zone: myHostTarget
00:3f:3f:3f:23:24:25:26; 3f:3f:3f:30:32:00:00:00
zone:red_0917_00_3f_3f_3f_23_24_25_26_3f_3f_3f_30_32_00_00_00
00:3f:3f:3f:23:24:25:26; 3f:3f:3f:30:32:00:00:00;
3f:3f:3f:30:30:00:00:00; 3f:3f:3f:30:31:00:00:00
zone: red_____base
00:00:00:00:00:00:00:01; 00:00:00:00:00:00:00:02;
00:00:00:00:00:00:00:03; 00:00:00:00:00:00:00:04
switch:admin> nszonestow wwn 00:3f:3f:3f:23:24:25:26

Zone Names
=====
myHostTarget
red_0917_00_3f_3f_3f_23_24_25_26_3f_3f_3f_30_32_00_00_00
=====
```

SEE ALSO [cfgShow](#), [zone](#), [nsZoneMember](#)

openSource

Displays open-source licenses.

Synopsis `opensource`

Description Displays listing of open source elements, licenses, and code samples used in Fabric OS. Refer to the Linux documentation for more information on how to use this command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display open-source licenses:

```
switch:user> opensource
In compliance with open source licensing requirements Brocade
publishes open source information in the following location:
http://www.brocade.com/support/oscd
Also included in this file is a list of all open source licences
associated with this release of the Fabric Operating System.
=====
License Text for apache 1.3.31
Relative path to license text file: LICENSE
=====
/* =====
* The Apache Software License, Version 1.1
*
* Copyright (c) 2000-2002 The Apache Software Foundation. All rights
* reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. The end-user documentation included with the redistribution,
* if any, must include the following acknowledgment:
* "This product includes software developed by the
* Apache Software Foundation (http://www.apache.org/)."
* Alternately, this acknowledgment may appear in the software itself,
* if and wherever such third-party acknowledgments normally appear.
*
Type <CR> or <SPACE BAR> to continue, <q> to stop
(output truncated)
```

SEE ALSO [aliShow](#)

passwd

Changes the password for a specified user.

Synopsis

```
passwd
  [user_account]
  [-old old_password]
  [-new new_password]
```

Description

Use this command to change a user account password.

Passwords can be changed locally on any switch. For the password database to be distributed to other switches in the fabric, the switches must be configured to accept the password database with the **fddCfg** command. The password database is distributed manually with the **distribute** command.

If RADIUS authentication is enabled, password change is blocked for users changing their own password. Administrators with the privilege to change passwords for other accounts may do so regardless of whether RADIUS authentication is enabled; all such password changes operate on the local password database.

The **passwd** command cannot be run on the Standby CP. When an admin account or a SecurityAdmin account changes the password for other accounts, it does not prompt for the current password, unless the target account is a factory or root account.

Any chosen password must satisfy the following password strength rules:

- Password contains the minimum required number of lowercase characters.
- Password contains the minimum required number of uppercase characters.
- Password contains the minimum required number of numeric characters
- Password contains the minimum required number of punctuation characters.
- Password must be between *minlength* and 40 characters long. The *minlength* parameter is set with the **passwdCfg** command.
- Password may not contain the colon (:) character.
- Password must satisfy repeated and sequential character constraints.

The password history policy is enforced across all user accounts when the user is setting his own password. The password history policy is not enforced when an administrator sets a password for another user, but the user's password history is preserved and the password set by the administrator is recorded in the user's password history.

The **passwd** command behaves as follows:

- If you are changing your own password, you are prompted to enter the old password and, if your entry is valid, you are prompted to enter the new password. Alternately, you may specify the old and new password on the command line.
- If you are changing another user's password with greater privileges than your current login level, you are prompted to enter that user level's old password and, if your entry is valid, you are prompted for a new password.
- If you are changing another user's password, the target account's AD member list must be a subset of your account's AD member list.
- If you are logged in as the root user when changing another user's password, you are not prompted to enter the old password. If you are a factory account, an admin, or SecurityAdmin account, you are not prompted to enter the current password unless the target account is root.
- Changing the password of any user level causes the login session of that account (if logged in) to terminate.

When invoked without an operand, this command changes the password for the current user account.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

The following operand is optional:

user_account

Specifies the user account for which the password is to be changed. The user account must be an existing account, either default or user-created. Only users with accounts of type "root", "factory", "SecurityAdmin", or "admin" can execute this operand and have permission to change passwords for accounts other than their own.

-old old_password

Specifies the old password. This operand is optional; if omitted, the command interactively prompts for the old password.

-new new_password

Specifies the new password. This operand is optional; if omitted and only the old password is specified on the command line, the command interactively prompts for the new password.

Diagnostics

The system may generate one or more of the following error messages. Refer to the *Fabric OS Message Reference* for more diagnostic information.

RADIUS enabled error

Password change disallowed, RADIUS authentication is enabled.

Password error

Password length, strength, or history policy not met.

"User" is not a valid user name

You have not specified a valid recognized user name on the system.

Permission denied

You do not have permission to change the password for the specified user.

Incorrect password

You have not entered the correct password when prompted for the old password.

Password unchanged

You have entered the carriage return special input case, choosing not to change the password.

Passwords do not match

You have not correctly verified the new password.

Invalid length of password

You have entered a password length that is not between minlength and 40.

Insufficient number of lower case letters

The password you entered contains less than the minimum required number of lowercase characters.

Insufficient number of upper case letters

The password you entered contains less than the minimum required number of uppercase characters.

Insufficient number of digits in password

The password you entered contains less than the minimum required number of numeric characters.

Insufficient number of punctuation characters

You have entered a password that contains less than the minimum required number of punctuation characters.

Password matches one of the previous passwords

The password you entered matches one of the previous passwords.

You must wait longer to change your password

You cannot change the password before the minimum aging period expires.

Password contains invalid characters

The password you entered contains invalid characters.

Examples

To change the password for the admin account while logged in as admin:

```
switch:admin> passwd
Changing password for admin
Enter new password:
Re-type new password:
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
```

To change the password for user "admin" from an admin account noninteractively:

```
switch:admin> passwd -old password -new adminpass
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
```

passwd

To change the password for user "brocadeUser" from an admin account noninteractively:

```
switch:admin> passwd brocadeUser -old brcdPasswd -new mynewPasswd  
Password changed.  
Saving password to stable storage.  
Password saved to stable storage successfully.
```

SEE ALSO [login](#), [logout](#), [passwdCfg](#)

passwdCfg

Manages the password policies.

Synopsis

```
passwdcfg --set options value
passwdcfg --setuser username options value
passwdcfg --disableadminlockout
passwdcfg --enableadminlockout
passwdcfg --setdefault
passwdcfg --showall
passwdcfg --showuser username
passwdcfg --deleteuser username
passwdcfg --deleteuser -all
passwdcfg --hash md5 | sha256 | sha512 [-manual]
passwdcfg --showhash [username | -all]
passwdcfg --help
```

Description

Use this command to manage password policies.

Use **--set** to configure the following password policies:

- Password strength policy
- Password history policy
- Password expiration policy
- Account lockout policy

Password Strength Policy

The *password strength policy* enforces a set of rules that new passwords must satisfy. Configurable rules include lowercase and uppercase characters, numbers, punctuation occurrences and minimum length values. It is enforced only when a new password is defined. The password strength policy is enforced cross all user accounts. When a password fails more than one of the strength attributes, an error is reported for only one of the attributes at a time.

Password History Policy

The *password history policy* prevents reuse of a recently used password. The password history policy is enforced across all user accounts when users are setting their own password. It is not enforced when an administrator sets a password for another user, but the user's password history is preserved and the password set by the administrator is recorded in the user's password history.

Password Expiration Policy

The *password expiration policy* forces expiration of a password after a specified period of time. When a user's password expires, the user must change the password to complete the authentication process. A warning that password expiration is approaching is displayed when the user logs in. The number of days prior to password expiration during which warnings commence is a configurable parameter. Password expiration does not disable or lock out the account. The password expiration policy is enforced across all user accounts except the root and factory accounts.

Account Lockout Policy

The *account lockout policy* disables a user account when the user exceeds a configurable number of failed login attempts. The mechanism can be configured to keep the account locked until explicit administrative action is taken to unlock the account or locked accounts can be automatically unlocked after a specified period. An administrator can unlock a locked account at any time. Note that the *account locked* state is distinct from the *account disabled* state. The account lockout policy is enforced across all user accounts except the root, factory, and SecurityAdmin role accounts. A separate configuration option, available to the SecurityAdmin and Admin role accounts, may be used to enable and disable applications of the account lockout policy to Admin role accounts.

A failed login attempt counter is maintained for each user on each switch instance. The counters for all user accounts are reset to zero when the account lockout policy is enabled. The counter for an individual account is reset to zero when the account is unlocked after the lock-out duration period expires.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command takes as input an operand and its associated arguments. When invoked without operands, the command prints the usage.

--showall

Displays the password configuration parameters.

--showuser *username*

Displays the password configuration parameters for a specific user.

--setdefault

Resets all password policies to their default values.

--set

Configures a specified password policy.

-charset *value*

Specifies the minimum criteria on the character set (upper and lowercase letters and special characters) to be used in the password. The default value is 0.

-allowuser [yes|no]

Activates (no) or deactivates (yes) the validation check to determine if the username is used in the password. If activated, the username both in forward and reverse direction cannot be used in the password. For example, if username is "testuser", the password "testuser123" or "resutset567" is not allowed. The default value is "yes".

-lowercase *value*

Specifies the minimum number of lowercase alphabetic characters that must occur in the password. The default value is 0. The maximum value must be less than or equal to the *-minlength value*.

-uppercase value

Specifies the minimum number of uppercase alphabetic characters that must occur in the password. The default value is 0. The maximum value must be less than or equal to the **-minlength value**.

-digits value

Specifies the minimum number of numeric digits that must occur in the password. The default value is 0. The maximum value must be less than or equal to the **-minlength value**.

-punctuation value

Specifies the minimum number of punctuation characters that must occur in the password. All displayable, non-alphanumeric punctuation characters, except the colon (:), are allowed. The default value is 0. The maximum value must be less than or equal to the **-minlength value**.

-minlength value

Specifies the minimum length of the password. The minimum can be set anywhere between 8 and 40 characters. The default value is 8. The total of **-lowercase**, **-uppercase**, **-digits**, **-punctuation** must be less than or equal to **-minlength value**.

-history value

Specifies the number of past password values that are disallowed when setting a new password. A value of 0 to 24 can be specified. The default value is 0.

-minpasswordage value

Specifies the minimum number of days that must elapse before a password can be changed. **-minpasswordage** can be set at 0 to 999. The default value is 0. Setting this parameter to a nonzero value discourages a user from rapidly changing a password in order to defeat the password history setting to reuse a recently used password. The **minpasswordage** policy is not enforced when an administrator changes the password for another user.

-maxpasswordage value

Specifies the maximum number of days that can elapse before a password must be changed. This is the password expiration period. **-maxpasswordage** can be set at 0 to 999. Setting this parameter to 0 disables password expiration. The default value is 0. When **-maxpasswordage** is set to a nonzero value, **-minpasswordage** must be set to a value less than or equal to **-maxpasswordage**.

-warning value

Specifies the number of days prior to password expiration that a warning of password expiration is displayed. The valid range for **-warning** is 0 to 999. The default value to 0.

-lockoutthreshold value

Specifies the number of times a user can specify an incorrect password during login before the account is locked. The number of failed login attempts is counted from the last successful login. Values for **-lockoutthreshold** range from 0 to 999. Setting this parameter to 0 disables the lockout mechanism. The default value is 0.

- lockoutduration *value***
Specifies the time, in minutes, after which a previously locked account automatically unlocks. **lockoutduration** values range from 0 to 99999. The default value is 30. Setting this parameter to 0 disables lockout duration, requiring an administrative action to unlock the account. The lockout duration begins with the first login attempt after the lockout threshold has been reached. Subsequent failed login attempts do not extend the lockout period.
- repeat *value***
Specifies the length of repeated character sequences that will be disallowed. For example, if the "repeat" value is set to 3, a password "passAAAword" is disallowed because it contains the repeated sequence "AAA". A password of "passAAword" would be allowed because no repeated character sequence exceeds two characters. The range of allowed values is 1 to 40.
- sequence *value***
Specifies the length of sequential character sequences that will be disallowed. In a character sequence, the ASCII value of each contiguous character differs by one. The ASCII value for the characters in the sequence must all be increasing or decreasing. For example, if the "sequence" value is set to 3, a password "passABCword" is disallowed because it contains the sequence "ABC". A password of "passABword" would be allowed because no repeated character sequence exceeds two characters. The range of allowed values is 1 to 40. The default value is 1.
- reverse [1|0]**
Activates (1) or deactivates (0) the validation check to determine whether the password is an exact reverse string of the username.
- oldpasswd [1|0]**
Enables or disables old password check while changing password for the root account using the **passwd** command.
- setuser *username***
Configures the password policy for a specific user.
- minpasswordage *value***
Specifies the minimum number of days that must elapse before a password can be changed. **-minpasswordage** can be set at 0 to 999. The default value is 0. Setting this parameter to a nonzero value discourages a user from rapidly changing a password in order to defeat the password history setting to reuse a recently used password. The **minpasswordage** policy is not enforced when an administrator changes the password for another user.
- maxpasswordage *value***
Specifies the maximum number of days that can elapse before a password must be changed. This is the password expiration period. **-maxpasswordage** can be set at 0 to 999. Setting this parameter to 0 disables password expiration. The default value is 0. When **-maxpasswordage** is set to a nonzero value, **-minpasswordage** must be set to a value less than or equal to **-maxpasswordage**.
- warning *value***
Specifies the number of days prior to password expiration that a warning of password expiration is displayed. The valid range for **-warning** is 0 to 999. The default value to 0.

- enableadminlockout**
Enables the admin lockout policy and sets the config parameter "passwdcfg.adminlockout" to 1. If the parameter "passwdcfg.lockoutthreshold" is set to greater than 0 and Admin Lockout policy is enabled, then, if the number of failed login attempts from the last successful login equals the "passwdcfg.lockoutthreshold", the account gets locked for the "passwdcfg.lockoutduration" duration. The particular account is unlocked manually using **userconfig --change account_name -u** or it is automatically unlocked after "passwdcfg.lockoutduration" duration.
- disableadminlockout**
Disables the admin lockout policy if already enabled and sets the config parameter "passwdcfg.adminlockout" to 0. By default, admin lockout policy is disabled.
- deleteuser username**
Removes the password expiration policies for a particular user account.
- deleteuser -all**
Removes the password expiration policies of all users.
- hash md5 | sha256 | sha512**
Sets the hash type. Valid values are "md5", "sha256", or "sha512". After password hash type is configured on switch, you will be prompted to change the passwords on the next login if hash type of password is different from the configured hash type on switch.
- manual**
Requires manual password change.
- showhash [username | -all]**
Displays password hash configuration for the given user or all users. When executed without optional parameters, this command displays the password hash configured in the system.
- help**
Displays the command usage.

Diagnostics

The **passwdCfg** command may fail for any of the following reasons:

Permission failure

You not permitted to execute the command.

Invalid command line option

An unrecognized command line option was specified.

Minlength value out of range

The **-minlength value** must be between 8 and 40.

Lowercase value out of range

The **-lowercase value** specified must be greater than or equal to 0 and less than or equal to **-minlength value**.

Uppercase value out of range

The **-uppercase value** specified must be greater than or equal to 0 and less than or equal to **-minlength value**.

Digits value out of range

The **-digits value** specified must be greater than or equal to 0 and less than or equal to **-minlength value**.

Punctuation value out of range

The **-punctuation value** specified must be greater than or equal to 0 and less than or equal to **-minlength value**.

Total strength specification out of range

The total of **-lowercase value**, **-uppercase value**, **-digits value**, and **-punctuation value** must be less than or equal to **-minlength value**.

History value out of range

The **-history value** must be between 0 and 24.

Minpasswordage value out of range

The **-minpasswordage value** must be between 0 and 999.

Maxpasswordage value out of range

The **-maxpasswordage value** must be between 0 and 999.

Warning value out of range

The **-warning value** must be between 0 and 999.

Invalid password expiration specification

The **-minpasswordage value** or **-warning value** must be less than or equal to **-maxpasswordage value** when **-maxpasswordage value** is nonzero.

Lockoutthreshold value out of range

The **-lockoutthreshold value** must be between 0 and 999.

Lockoutduration value out of range

The **-lockoutduration value** must be between 0 and 99999.

Repeat value out of range

The **-repeat value** must be between 1 and 40.

Sequence value out of range

The **-sequence value** must be between 1 and 40.

Specified user account not exist

The user does not exist in the local switch.

Per user password expiration policy not configured

The password expiration policy is not configured for this user.

Not applicable for default account

The password expiration policy is not applicable for the root and factory accounts.

Password strength policy is allowed for per user password configurations

The password configuration is allowed only for the password expiration policy.

Examples To display the current password configuration parameters:

```
switch:admin> passwdcfg --showall
passwdcfg.minlength: 8
passwdcfg.lowercase: 0
passwdcfg.uppercase: 0
passwdcfg.digits: 0
passwdcfg.punctuation: 0
passwdcfg.history: 1
passwdcfg.minpasswordage: 0
passwdcfg.maxpasswordage: 0
passwdcfg.warning: 0
passwdcfg.lockoutthreshold: 0
passwdcfg.lockoutduration: 30
passwdcfg.status: 0
passwdcfg.reverse: 0
```

To set password configuration parameters, specifying that a password must contain at least two uppercase characters, and that passwords expire in 90 days from the date they are defined:

```
switch:admin> passwdcfg --set -uppercase 2 \
-maxpasswordage 90
```

To set user password expiration policy parameters:

```
switch:admin> passwdcfg --setuser admin -minpasswordage 2 \
-maxpasswordage 5 -warning 3
```

To display the current user password expiration policy parameters:

```
switch:admin> passwdcfg --showuser admin
passwdcfg.maxpasswordage.admin: 5
passwdcfg.minpasswordage.admin: 2
passwdcfg.warning.admin: 3
```

To delete all password configurations:

```
switch:admin> passwdcfg --deleteuser -all
All user password configurations are removed
```

To delete password configurations of a specific user:

```
switch:admin> passwdcfg --deleteuser admin
The user(admin) password configurations are removed
```

To activate password reverse criterion check:

```
switch:admin> passwdcfg --set -reverse 1
```

SEE ALSO [passwd](#), [userConfig](#)

pathInfo

Displays routing information and statistics along a path covering multiple switches.

Synopsis

pathinfo

```
pathinfo -f FID dest_switch_domain [source_port] [dest_port]
          [-sid source_pid] [-did dest_pid] [-r] [-t]
```

```
pathinfo dest_switch_domain [source_port] [dest_port]
          [-sid source_pid] [-did dest_pid] [-r] [-t]
```

```
pathinfo --help
```

Description

Use this command to display routing information from a source port on the local switch to a destination port on another switch. The command output describes the exact data path between these ports, including all intermediate switches.

When using **pathInfo** across fabrics connected through an FC Router, the command represents backbone information as a single hop. The command captures details about the FC Router to which ingress and egress EX_Ports are connected, but it hides the details about the path the frame traverses between the ingress EX_Ports to the egress EX_Ports in the backbone.

To use **pathInfo** across remote fabrics, you must specify both the fabric ID (FID) and the domain ID of the remote switch. You cannot use the command to obtain source port information across remote FCR fabrics. When obtaining path info across remote fabrics, the destination switch must be identified by its Domain ID. Identifying the switch by name or WWN is not accepted.

The command does not retry if there is a timeout or failure. It may fail if a switch along the path is busy or does not support this feature.

If the advanced performance tuning (APT) policy in effect on the intermediate switches is not a port-based policy, subsequent data streams may not take the same path as displayed in the **pathInfo** output. Refer to **aptPolicy** for more information on advanced performance tuning policies.

If you specify an inactive port or a path through a switch that does not have active routing tables to the destination, this command describes the path that would be taken if the ports were active. If you specify a destination port that is not active, this command uses the embedded port as the destination.

The ingress and egress points above 256 are specified as port index. Use **switchShow** for a listing of valid port index numbers.

To display the TI-enabled path across FCR, you must specify both source ID (SID) and destination ID (DID). The DID must be the proxy device ID and you can view the proxy device IDs using the **nsAllShow** command.

In addition, **pathInfo** can provide statistics on every traversed interswitch link (ISL) that is part of the path. This feature is available only in the interactive command mode.

The routing and statistics information are provided by every switch along the path, based on the current routing table information and statistics calculated continuously in real-time. Each switch represents one hop of the total path.

In a Virtual Fabric environment, **pathInfo** displays port numbers beyond physical port numbers while displaying information for a hop that corresponds to a path in the base fabric. The cost for this hop is the cost of the corresponding path in the base fabric plus a small offset. Refer to the Examples section for an illustration.

Other command options allow the collection of information on the reverse path, or on a user-selected path (source route).

For each hop, this command displays the following fields:

| | |
|------------------|--|
| Hop | The hop number. The local switch is hop 0. |
| In Port | The port on which the switch receives frames. For hop 0, this is <i>source_port</i> , identified by the port index. |
| Domain ID | The domain ID of the switch. |
| Name | The name of the switch. |
| Out Port | The output port that the frames take to reach the next hop. For the last hop, this is <i>destination_port</i> identified by the port index. |
| BW | The bandwidth of the output interswitch link (ISL), in Gbps. This parameter does not apply to the embedded port. If the bandwidth is zero, it is displayed as 1 Gbps. For logical interswitch links (LISL) ports, the bandwidth displays as 8 Gbps, the maximum bandwidth on hardware platforms on which LISLs can be formed. In cases where the LISL bandwidth is zero, pathInfo displays a bandwidth of 4 Gbps. |
| Cost | The cost of the output link used by the fabric shortest path first (FSPF) routing protocol. This parameter applies only if the output link is recognized by FSPF. |

You can request more detailed statistics for each hop in addition to the routing information. These statistics are presented for the input and output ports for both receive and transmit modes. You can select basic or extended statistics or both when running **pathInfo** in interactive mode. Statistics are not reported for the embedded port. Some throughput values are reported in multiple time intervals, to describe both current path utilization and the average throughput over a larger period of time.

To collect these statistics, this command uses a special **pathInfo** frame that is sent hop-by-hop from the source switch to the destination switch. To prevent such a frame to loop forever if an error occurs, a maximum number of hops for the frame to traverse is enforced. The hop count includes all hops in the direct path from source to destination, and also all the hops in the reverse path, if the tracing of the reverse path is requested. The default value for the maximum hop count is 25.

Basic statistics

Basic statistics report parameters that may indicate ISL congestion along the path. They include the following:

B/s

Bytes per second received or transmitted. This value is reported for multiple time periods displayed in parentheses.

Txcrdz

The length of time, in milliseconds, that the port was unable to transmit frames because the transmit BB credit was zero. The purpose of this statistic is to detect congestion or a device affected by latency. This parameter is sampled at 1 millisecond intervals, and the counter is incremented if the condition is true. Each sample represents 1 ms of time with zero Tx BB Credit. An increment of this counter means that the frames could not be sent to the attached device for 1 ms, indicating degraded performance. This value reports for multiple time periods, displayed in parentheses. Note that other commands, such as **portStatsShow**, may express this value in units other than milliseconds.

Extended statistics

Extended statistics report variables of general interest. They include the following:

F/s

The number of frames received or transmitted per second. This value is reported for multiple time periods, displayed in parentheses.

Words

The total number of 4-byte Fibre Channel words.

Frames

The total number of frames.

Errors

The total number of errors that may have caused a frame not to be received correctly. This includes cyclic redundancy check (CRC) errors, bad end-of-frame (EOF) errors, frame truncated errors, frame-too-short errors, and encoding errors inside a frame.

Reverse path

The path from port A on switch X to port B on switch Y may be different from the path from port B to port A depending on the links traversed between a given sequence of switches, or the reverse path may involve different switches. The **-r** option displays routing and statistics information for the reverse path in addition to those for the direct path.

Source route

The source route option allows you to specify a sequence of switches or ports, which the **pathInfo** frame has to traverse to reach the destination. Therefore, the path specified may be different from the one used by actual traffic.

The source route is expressed as a sequence of switches, a sequence of output ports, or a combination of both. The next hop in the source route is described by either the output port to be used to reach the next hop, or the domain ID of the next hop.

The source route can specify a full route from source to destination or a partial route. In a partial route the remaining hops are chosen as the path from the input port on the first hop not listed in the source route to the destination. The maximum hop count is enforced in both cases.

If the source route does not specify all the switches along a section of the path, you can specify a strict or a loose path. A strict source route requires that only the specified switches are reported in the path description. If two switches are specified back-to-back in the source route descriptor but are not directly connected, the switches in-between are ignored. In a loose source route, the switches in-between are reported. The concepts of strict and loose route apply only to the portions of the path described by domains, not to the part described by output ports.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

dest_switch_domain

Specifies the destination switch. To obtain path info in a Layer 2 fabric, the destination switch can be identified by its Domain ID, by the switch WWN, or by the switch name. To obtain path info across remote fabrics connected through an FC Router, the destination switch must be identified by its Domain ID. Identifying the switch by name or WWN is not accepted. This operand is optional; if omitted, the command prompts for input interactively.

source_port

Specifies the port whose path to the destination domain is traced, specified as the port index. The embedded port (-1) is the default. The embedded port can be selected manually by entering the value of MAX_PORT. MAX_PORT stands for the maximum number of ports supported by the local switch.

dest_port

Specifies the port on the destination switch for the path being traced. This operand returns the state of this port. The embedded port (-1) is used by default, or if you specify a destination port that is not active. The destination is specified as the port index.

"-r"

Displays the reverse path in addition to the forward path. This operand is optional.

-t

Displays the command output in traceroute format. When this operand is used, only routing information is displayed. The output includes the time it takes, in microseconds, to reach each hop. Basic and extended statistics are not available in the traceroute format.

-sid source_pid

Specifies the source ID of the originator device in "OxDDAAPP" format, where DD is domain ID, AA is area ID and PP is AL_PA ID.

-did dest_pid

Specifies the destination ID of the proxy device in "OxDDAAPP" format, where DD is domain ID, AA is area ID and PP is AL_PA ID. You can view the proxy device IDs using the **nsAllShow** command.

pathInfo

--help

Displays the command usage.

When executed without operands, **pathInfo** runs interactively. Specifying a destination switch is required; the values for the source and destination ports can be -1 to indicate the embedded port. You must specify the source and destination PID to display the correct TI-enabled path information across the FCR. Reverse-path tracing is optional. In addition, this command prompts for the following parameters:

Max hops

The maximum number of hops that the **pathInfo** frame is allowed to traverse; the default is 25.

Fabric Id

Specifies the Fabric ID. If unspecified, the value defaults to -1 (Default switch FID 128)

Domain|Wwn|Name

Specifies the domain ID or WWN.

Source port

Specifies the port whose path to the destination domain is traced, specified as the port index. If unspecified, the value defaults to -1 (embedded port).

Destination port

Specifies the port on the destination switch for the path being traced. If unspecified, the value defaults to -1 (embedded port)

Source pid

Specifies the source ID of the originator device in hexadecimal format.

Destination pid

Specifies the destination ID of the proxy device in hexadecimal format.

Basic stats

Whether basic statistics are generated on every link; the default is no.

Extended stats

Whether extended statistics are generated on every link; the default is no.

Trace reverse path

Specifies the path information in reverse direction; the default is no.

Source route

Specifies a sequence of switches or ports that the **pathInfo** frame should traverse; the default is no. If an output port to the next hop is specified, you are not prompted for the domain of the next switch. The domain is determined by the port.

Timeout

The maximum time allowed waiting for the response in milliseconds. The default is 10000 milliseconds.

Examples To display basic path information to a specific domain in command line mode:

```
switch:admin> pathinfo 91
Target port is Embedded
Hop  In Port  Domain ID (Name)      Out Port  BW      Cost
-----
0      E        9 (web226)            2         1G     1000
1      3        10 (web229)           8         1G     1000
2      8        8 (web228)            9         1G     1000
3      6        91 (web225)           E         -      -
```

To display basic path information in traceroute format:

```
switch:admin> pathinfo 91 -t
traceroute to domain 91(web225) , 25 hops
Hop  Domain ID (Name)      time taken for the hop
-----
1    10 (web229)           0.0331 us
```

To display basic path information in traceroute format with reverse path option:

```
switch:admin> pathinfo 4 -r -t
Target port is Embedded
Hop      Domain ID (Name)      Time/hop
-----
1         11 (mps_daz_1)       32882 usec
2          4 (METEOR)         32882 usec
3         11 (mps_daz_1)       32882 usec
4         97 (pulsar055)       32882 usec
```

To display path information when source port and destination port are provided along with the traceroute option:

```
switch:admin> pathinfo 6 12 13 -t
Target port is F_Port
Hop      Domain ID (Name)      Time/hop
-----
6 (Stealth_I)           108186 usec
```

To display basic path information to a specific domain in a Virtual Fabric environment (the cost for this hop is the cost of the corresponding path in the base fabric):

```
switch:admin> pathinfo 13 4
Target port is Embedded
Hop  In Port  Domain ID (Name)      Out Port  BW      Cost
-----
0      4        9                    2009*    -      1500
1     2015*    10                    8         1G     1000
2      6        13                    E         -      -
```

pathInfo

To display basic and extended statistics in interactive mode:

```
switch:admin> pathinfo
Max hops: (1..127) [25]
Fabric Id: (1..128) [-1]
Domain|Wwn|Name: [] 8
Source port: (0..200) [-1]
Destination port: (0..1800) [-1]
Source pid: (0x0..0xffff00) [ffffffff]
Destination pid: (0x0..0xffff00) [ffffffff]
Basic stats (yes, y, no, n): [no] y
Extended stats (yes, y, no, n): [no] y
Trace reverse path (yes, y, no, n): [no]
Source route (yes, y, no, n): [no]
Timeout: (1000..30000) [5]
Target port is Embedded
```

| Hop | In Port | Domain ID (Name) | Out Port | BW | Cost |
|-----|---------|------------------|----------|----|------|
| 0 | E | 9 (web226) | 2 | 1G | 1000 |

| Port | E | | 2 | |
|--------------|----|----|---------|---------|
| | Tx | Rx | Tx | Rx |
| B/s (1s) | - | - | 0 | 0 |
| B/s (64s) | - | - | 1 | 1 |
| Txcrdz (1s) | - | - | 0 | - |
| Txcrdz (64s) | - | - | 0 | - |
| F/s (1s) | - | - | 0 | 0 |
| F/s (64s) | - | - | 2743 | 0 |
| Words | - | - | 2752748 | 2822763 |
| Frames | - | - | 219849 | 50881 |
| Errors | - | - | - | 0 |

| Hop | In Port | Domain ID (Name) | Out Port | BW | Cost |
|-----|---------|------------------|----------|----|------|
| 1 | 3 | 10 (web229) | 12 | 1G | 1000 |

| Port | 3 | | 12 | |
|--------------|-----------|---------|---------|---------|
| | Tx | Rx | Tx | Rx |
| B/s (1s) | 36 | 76 | 0 | 0 |
| B/s (64s) | 5 | - | 0 | - |
| Txcrdz (64s) | 0 | - | 0 | - |
| F/s (1s) | 1 | 1 | 0 | 0 |
| F/s (64s) | 0 | 0 | 0 | 0 |
| Words | 240434036 | 2294316 | 2119951 | 2121767 |
| Frames | 20025929 | 54999 | 162338 | 56710 |
| Errors | - | 4 | - | 0 |

| Hop | In Port | Domain ID (Name) | Out Pot | BW | Cost |
|-----|---------|------------------|---------|----|------|
| 2 | 14 | 8 (web228) | E | - | - |

(output truncated)

To display the TI-enabled path information over an FCR in interactive mode:

```

switch:admin> pathinfo
Max hops: (1..127) [25]
Fabric Id: (1..128) [-1] 8
Domain|Wwn|Name: [] 5
Source port: (0..400) [-1] 2
Destination port: (0..400) [-1] 24
Source pid: (0x0..0xffff00) [ffffff] 0x061600
Destination pid: (0x0..0xffff00) [ffffff] 0x01f001
Basic stats (yes, y, no, n): [no]
Extended stats (yes, y, no, n): [no]
Trace reverse path (yes, y, no, n): [no] y
Source route (yes, y, no, n): [no]
Timeout: (1000..30000) [10000]

Target port is Embedded

Hop   In Port   Domain ID (Name)      Out Pot   BW      Cost
-----
0     2         1 (web228)           6         4G      500
1     23        2 (web228)           8         4G      500
2     4         3 (web228)           3         4G      500
3     2         4 (web228)           24        4G      10000
4     3         7 (switch_3)         2         4G      500
5     27        5 (switch_3)         24        -       -

Reverse path
6     24        5 (switch_3)         27        4G      500
7     2         7 (switch_3)         3         4G      500
8     24        4 (web228)           2         4G      500
9     3         3 (web228)           4         4G      10000
10    8         2 (web228)           23        4G      500
11    6         1 (web228)           2         -       -

```

SEE ALSO [portStatsShow](#), [switchShow](#)

pdShow

Displays data from a panic dump file.

Synopsis `pdshow [panic_dump_file]`

Description Use this command to display data from a panic dump file. The panic dump file contains information that might be useful to determine the cause of the system panic.

When executed without any arguments, this command displays output from the latest panic dump file available on the switch.

If a panic dump file is specified as an argument, the contents of that specific file are displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operand:

panic_dump_file Specify the full path name of a panic dump file.

Examples To examine a panic dump file by the name *panic_dump* located under the directory */tmp*:

```
switch:admin> pdshow /tmp/panic_dump
*** CAUTION ***
* Host PLATFORM (current) is: 'Unknown'
* PLATFORM got from pd file is: 'SW12000'
* Some results shown may be incorrect and/or missing
* It is best if this command is run on same \
  PLATFORM as that of pdfile
*****
*****
* File    :/core_files/panic/core.pd1038932352  *
* SECTION:PD_MISC                               *
-----*****-----
WatchDogRegister=0x0
Section=Startup time: Tue Dec  3 16:06:11 UTC 2002
Kernel=      2.4.19
Fabric OS=   v4.1.0_j_dist_1103
Made on=    Tue Dec 3 19:07:13 2002
Flash=      Tue Dec 3 13:19:06 2002
BootProm=   3.2.0
Section=HA show Output

(output truncated)
```

SEE ALSO [portLogDump](#), [supportSave](#)

portAddress

Assigns the lower 16 bits of the Fibre Channel Port ID.

Synopsis `portaddress --bind [slot/]port [16-bit_address] [--auto]`

`portaddress --unbind [slot/]port`

`portaddress --show [[slot/]port]`

`portaddress --findPID 24-bit_Port_ID`

`portaddress --help`

Description

Use this command to bind the 16-bit address to the lower two bytes of a port 24-bit Fibre Channel address, or to unbind the currently bound address for the specified port. Changes effected by this command are persistent across reboots and power cycles.

The port must be offline to bind an address and not currently bound to another address. If the port is currently bound to another address, use this command with the **--unbind** option to unbind the port.

This command returns an error if the chosen address is in use or is bound to another port. If the address is currently assigned to another port, use this command with the **--findPID** option to identify the port that is bound to that address, and then unbind the port.

The command provides a **--show** option that displays the currently bound address for a specified port or for all ports. Alternately, you can use the **--findPID** option to display the port currently bound to a specified port ID (PID).

Notes

This command is supported on the Brocade DCX 8510-8 and DCX 8510-4 on all logical switches including the Default Switch. Virtual Fabrics must be enabled. It is also supported on the Brocade 6510, regardless of Virtual Fabrics status.

This command is not supported on embedded platforms.

If a blade port is qualified to be in a user-created logical switch on a Brocade DCX 8510-4 or DCX 8510-8, this feature is supported on those ports. In the Default Switch, AP blades such as the Brocade FCOE10-24 or FS8-18 are not supported, when dynamic area mode is enabled. This is enforced by the software.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

--bind

Assigns the lower two bytes of the Fibre Channel address to the specified port.

slot

Specifies the slot number on bladed systems, followed by a slash (/).

port

Specifies the port number, relative to its slot on bladed systems.

16-bit_address

Specifies the 16-bit address, in hexadecimal format, to be bound to the FC address. Leading zeros are optional in the hexadecimal value (for example, you can either specify 0x1a00 or 1a00). Note that only the upper 10 bits of the PID can be used for a unique route. Therefore, not all addresses in the 16-bit range are available.

--auto

Enables autobinding on the specified port. If the auto feature is enabled, the entire area field of the PID is bound to a single port. With 10-bit routing, up to 4 ports can share the same 8-bit area field of the PID. This address assignment mode dedicates all four unique routes to a single port. By default, auto is off. This operand is optional; if unspecified, the default is used.

--unbind

Removes both the address and any automode override configuration from the specified port.

--show

Displays the currently bound address attributes for the specified port. This command shows the lowest two bytes of the Fibre Channel address as well as the current setting for automode. If a port is not specified, the display shows the Partition Address Mode value (0, 1, or 2) and all ports on the current partition. A -1 is displayed for ports that have not been assigned an area. Areas are dynamically assigned an address as they are added to a partition. The Partition Address Mode value is set by the **configure** command (Enable a 256 Area Limit).

--findPID

Displays the port (slot and port offset) of the port that is currently assigned the provided address. This command applies the 10-bit area mask to the provided PID and returns the port that has been assigned the specified address. Therefore not all 24 bits are required to match exactly.

24-bit_Port_ID

Specifies the 24-bit Fibre Channel port address. This operand is required with the **--findPID** option. This command applies the 10-bit area mask to the provided PID and returns the port that has been assigned the specified address. Therefore not all 24 bits are required to match.

--help

Displays the command usage.

Examples

To bind a 16-bit address to the low two bytes of a port 24-bit Fibre Channel address:

```
switch:admin> portaddress --bind 5/18 1a00
```

To unbind a given address from a port:

```
switch:admin> portaddress --unbind 5/18
```

To display all port address bindings on the current partition:

```
switch:admin> portaddress --show
Partition Address Mode :0
Index Slot Port Area Mode User_bound
=====
384 5 0 0x0800 8 bit Y
385 5 1 0x0900 8 bit -
386 5 2 0x0a00 8 bit -
387 5 3 0x0b00 8 bit -
388 5 4 0x0c00 8 bit -
389 5 5 0x0d00 8 bit -
390 5 6 0x0e00 8 bit -
391 5 7 0x0f00 8 bit -
392 5 8 0x0000 8 bit -
393 5 9 0x0100 8 bit Y
394 5 10 0x0200 8 bit -
395 5 11 0x0300 8 bit -
396 5 12 0x0400 8 bit -
397 5 13 0x0500 8 bit -
398 5 14 0x0600 8 bit -
399 5 15 0x0700 8 bit -
400 5 16 0x1800 8 bit -
401 5 17 0x1900 8 bit -
402 5 18 0x1a00 8 bit -
403 5 19 0x1b00 8 bit -
404 5 20 0x1c00 8 bit -
405 5 21 0x1d00 8 bit -
406 5 22 0x1e00 8 bit -
407 5 23 0x1f00 8 bit -
408 5 24 0x1000 8 bit -
409 5 25 0x1100 8 bit -
410 5 26 0x1200 8 bit -
411 5 27 0x1300 8 bit -
412 5 28 0x1400 8 bit Y
413 5 29 0x1500 8 bit -
414 5 30 0x1600 8 bit -
415 5 31 0x1700 8 bit -
```

To display the port address binding for port 28:

```
switch:admin> portaddress --show 5/18

Index Slot Port Area Mode
=====
412 5 28 0x1400 8 bit
```

To display the port bound to a specified address.

```
switch:admin> portaddress --findPID 0x2400

Index Port Port ID
=====
36 36 0x 2400
```

SEE ALSO None

portAlpaShow

Displays the Arbitrated Loop Physical Addresses (AL_PAs) of devices attached to a port.

Synopsis `portalpashow [slot/]port`

Description Use this command to display the AL_PAs of devices connected to a port, and whether these devices are public or private. If the specified port is not an active FL_Port or if no AL_PAs are present, this command prints an error.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot For bladed systems only, specify the slot number of the port to display, followed by a slash (/).

port Specify the port number to display, relative to its slot for bladed systems. Use **switchShow** to list valid ports.

Examples To display the AL_PAs of a port:

```
switch:user> portalpashow 4/14
AL_PA  type      AL_PA  type      AL_PA  type
0xe2   public      0xe4   public
```

SEE ALSO [switchShow](#)

portBeacon

Sets port beaconing mode.

Synopsis `portbeacon --enable [slot/]port`

`portbeacon --disable [slot/]port`

`portbeacon --show [slot/]port`

`portbeacon --show -all`

`portbeacon --help`

Description Use this command to enable or disable beaconing mode on a specified port.

When beaconing mode is enabled on a port, the port LED flashes amber and green for 2.5 seconds each in an alternating pattern. The beaconing mode continues until you turn it off. Beaconing mode is useful if you are trying to locate a specific port.

Beaconing mode takes over the port LEDs. The normal flashing LED pattern associated with an active, faulty, or disabled port is suppressed, and only the beaconing pattern is shown. Other commands are still executable and functional. However, if diagnostic frame-based tests such as **portLoopbackTest** are executed, the diagnostic LED pattern is interleaved with the beaconing pattern. Running switch beaconing or HBA- side (E2E) beaconing also overwrites the pattern.

The **portBeacon** command is one of the commands that controls beaconing. Each command has a clearly defined scope of action:

- The **portBeacon** command enables or disables beaconing on a specified port.
- The **switchBeacon** command enables or disables beaconing on all ports in the current logical switch.
- The **chassisBeacon** command enables or disables beaconing on all ports in the chassis.
- The **portPeerBeacon** command enables or disables beaconing to identify the interconnections between ports.

The actions of the beaconing commands are independent and mutually exclusive. Existing **portBeacon** settings remain unaffected if you enable or disable beaconing on the switch or on the chassis. Failure to disable existing beaconing commands before using a different type of beaconing may cause the commands to interfere with each other in unexpected ways.

Issue the **portBeacon --show [slot/]port** command to display beaconing for a specific port. The **switchShow** command displays the status of the **switchBeacon** command only.

Notes Beaconing of Inter Chassis Link (ICL) ports, GE ports, FCoE ports, 10G ports, and embedded internal ports is not allowed.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|--------------------|---|
| <i>slot</i> | For bladed systems only, specifies the slot number of the port group to display, followed by a slash (/). |
| <i>port</i> | Specifies the number of the port to be configured, relative to its slot for bladed systems. Use switchShow for a listing of valid ports. |
| --enable | Enables beaconing mode on the specified port. |
| --disable | Disables beaconing mode on the specified port. |
| --show | Displays the port beaconing mode on the specified port as ON or OFF. |
| --show -all | Displays the beaconing ports in slot/port format on the chassis and index format on switches, in logical switch order. |
| --help | Displays the command usage. |

Examples To enable beaconing mode on a port and verify the configuration:

```
switch:admin> portbeacon --enable 2/1

switch:admin> portbeacon --show 2/1
PortBeacon status of : slot 2, port 1 is ON
```

To disable beaconing mode on a port and verify the configuration:

```
switch:admin> portbeacon --disable 2/1

switch:admin> portbeacon --show 2/1
PortBeacon status of : slot 2, port 1 is OFF
```

To display the beaconing ports on a chassis:

```
switch:admin> portbeacon --show -all
In FID: 128
PortBeacon enabled port(s):
12/23 12/26
```

SEE ALSO [chassisBeacon](#), [portPeerBeacon](#), [switchBeacon](#)

portBufferCalc

Calculates the number of buffers required per port.

Synopsis `portbuffercalc [slot]port [-distance distance]
[-speed speed] [-framesize framesize]`

`portBufferCalc --help`

Description Use this command to calculate how many buffers are required for a given distance, speed, and framesize. You can use the same number of buffers and port speed, when configuring the **portCfgLongDistance** command. If no options are specified, then the current port's configuration is considered to calculate the number of buffers required.

Notes Buffers required for the given frame size, distance and speed is calculated using the below formula:

- 1 buffer is required for 1 km at 2G (2048 framesize).
- 8 buffers are required for 1 km at 16G (2048 framesize).

Operands This command supports the following operands:

| | |
|------------------------------------|---|
| -distance <i>distance</i> | Specifies the desired distance with valid range of 10Km to 1500Km. |
| -speed <i>speed</i> | Specifies the port speed. The valid values are: |
| 1 | Specifies port speed as 1Gbps. |
| 2 | Specifies port speed as 2Gbps. |
| 4 | Specifies port speed as 4Gbps. |
| 8 | Specifies port speed as 8Gbps. |
| 10 | Specifies port speed as 10Gbps. |
| 16 | Specifies port speed as 16Gbps. |
| 32 | Specifies port speed as 32Gbps. |
| -framesize <i>framesize</i> | Specifies the desired framesize in bytes with a maximum value 2048. |
| --help | Displays command usage. |

portBufferCalc

Examples To calculate the buffers required per port:

```
switch:admin> portbuffercalc 9/4 \  
-distance 100 -speed 8 -framesize 512  
1606 buffers required for 100km at 8G and framesize  
of 512 bytes
```

SEE ALSO [portCfgLongDistance](#)

portBufferShow

Displays the buffer usage information for a port group or for all port groups in the switch.

Synopsis `portbuffershow [[slot]port]`

Description Use this command to display the current long distance buffer information for the ports in a port group. The port group can be specified by giving any port number in that group. If no port is specified, then the long distance buffer information for all of the port groups of the switch is displayed.

The following long distance information is displayed:

| | |
|---|--|
| User Port | Index number of the port. |
| Port Type | E (E_Port), F (F_Port), G (G_Port), L (L_Port), or U (U_Port). |
| Lx Mode | Long distance mode. |
| L0 | Link is not in long distance mode. |
| LE | Link is up to 10 km. |
| LD | Distance is determined dynamically. |
| LS | Distance is determined statically by user input. |
| Max/Resv Buffers | The maximum or reserved number of buffers that are allocated to the port based on the estimated distance (as defined by the <i>desired_distance</i> operand of the <code>portCfgLongDistance</code> command). If the port is not configured in long distance mode, certain systems might reserve buffers for the port. This field then displays the number of buffers reserved for the port. |
| Avg Buffer Usage & FrameSize | The average buffer usage and average frame size for Tx and Rx. |
| Buffer Usage | The actual number of buffers allocated to the port. In LD mode, the number is determined by the actual distance and the user-specified desired distance (as defined by the <i>desired_distance</i> operand of the <code>portCfgLongDistance</code> command). |
| Needed Buffers | The number of buffers needed to utilize the port at full bandwidth (depending on the port configuration). If the number of Buffer Usage is less than the number of Needed Buffers , the port is operating in the buffer limited mode. |

Link Distance

For LO (not in long distance mode), the command displays the fixed distance based on port speed, for instance: 10 km (1 Gbps), 5 km (2 Gbps), 2 km (4 Gbps), or 1 km (8 Gbps). For static long distance mode (LE), the fixed distance of 10 km displays. For LD mode, the distance in kilometers displays as measured by timing the return trip of a MARK primitive that is sent and then echoed back to the switch. LD mode supports distances up to 500 km. Distance measurement on a link longer than 500 km might not be accurate. If the connecting port does not support LD mode, it shows "N/A".

Remaining Buffers

The remaining (unallocated and reserved) buffers in a port group.

A hyphen in one of the display fields indicates that no relevant information is available; there may be no connection to a port, or the port is disabled, or the port is not an E_Port.

When invoked without operands, this command displays the long distance buffer information for all the port groups of the switch.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

The following operands are optional:

slot

For bladed systems only, specifies the slot number of the port group to display, followed by a slash (/).

port

Specifies the number of a port associated with the port group, relative to its slot for bladed systems. Use **switchShow** for a list of valid ports.

Examples

To display the port buffer information:

```
switch:admin> portbuffershow 17
User Port Lx Max/Resv Avg Buffer Usage & FrameSize\
Port Type Mode Buffers Tx Rx \
-----
Buffer Needed Link Remaining
Usage Buffers Distance Buffers
-----
64 - 8 - ( - ) - ( - )\
0 - -
65 - 8 - ( - ) - ( - )\
0 - -
66 - 8 - ( - ) - ( - )\
0 - -
67 - 8 - ( - ) - ( - )\
0 - -
68 E LS 806 197(2012) 201(2044)\
206 206 100km
69 E - 8 1(2016) 1(2020)
26 26 2km
70 E - 8 1(2012) 1(2036)\
26 26 2km
71 E - 8 1(2008) 2(2052)\
26 26 2km
```

```

192      -      8      - ( - )      - ( - )\
0      -      -
193      -      8      - ( - )      - ( - )\
0      -      -
194      -      8      - ( - )      - ( - )
0      -      -
195      -      8      - ( - )      - ( - )\
0      -      -
196      -      8      - ( - )      - ( - )\
0      -      -
197      -      8      - ( - )      - ( - )\
0      -      -
198      -      8      - ( - )      - ( - )\
0      -      -
199      -      8      - ( - )      - ( - )\
0      -      -
-----
4556

```

SEE ALSO [portCfgLongDistance](#)

portCamShow

Displays port-based filter CAM utilization.

Synopsis `portcamshow [slot/]port`

Description Use this command to display the current filter Content-Addressable Memory (CAM) utilization of a specified port.

The command displays the following information:

SID used

Total number of CAM entries used by this port. Note that each CAM entry, either source ID (SID) or destination ID (DID) CAM, can be shared among a certain number of ports, depending on the ASIC.

DID used

Total number of CAM entries used by this port. Note that each CAM entry (either SID or DID CAM) can be shared among a certain number of ports, depending on the ASIC.

SID entries

All existing source ID entries within the CAM for this port. Note that each CAM entry (either SID or DID CAM) can be shared among a certain number of ports, depending on the ASIC.

DID entries

All existing destination ID entries within the CAM for this port. Note that each CAM entry (either SID or DID CAM) can be among a certain number of ports, depending on the ASIC.

SID free

The total number of free SID CAM entries available for use by this port.

DID free

The total number of free DID CAM entries available for use by this port.

Notes This command cannot be executed on a logical port.

Ports that support shared areas are divided into two ports: primary and secondary ports. Primary and secondary ports share the same area. Port CAM entries displayed on the primary ports also consist of all the secondary port SIDs and DIDs as well, when both the ports are F_Ports. This is because the primary port acts as a proxy for the CAM entries of the secondary port, in this case using redirect filters.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports.

This command cannot be executed on a logical port.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to display, followed by a slash (/).

port

Specifies the port number to display, relative to its slot for bladed systems. Use **switchShow** to list valid ports. This operand is required.

Examples

To display the filter CAM utilization for a single port on a switch:

```
switch:user> portcamshow 3/2
-----
Area  SID used  DID used  SID entries  DID entries
 34    3        1      350400      2b2200
                2b1200
                220400
-----
SID Free, DID Free: (61, 511)
```

To display port CAM entries on shared ports (in the example, port 7/31 and 7/39 are shared ports and 7/31 is the primary port):

```
switch:user> portcamshow 7/39
-----
Area  SID used  DID used  SID entries  DID entries
207    3        1      03b380      03cf80
                034100
                03cf00
-----
SID free, DID free: (2044, 1020)

switch:admin> portcamshow 7/31
-----
Area  SID used  DID used  SID entries  DID entries
207    4        2      03b380      03cf80
                034100      03cf00
                03cf00
                03cf80
-----
SID free, DID free: (2044, 1020)
```

The SID entry 03cf00 and DID entry 03cf80 on port 7/31 belong to port 7/39.

SEE ALSO [switchShow](#)

portCfg

Manages port configuration parameters for FC ports, VE_ports, and GbE/10GbE ports.

Synopsis

portcfg action [*slot*] *port arguments*

portcfg action [*slot*] *ge_port arguments*

portcfg action [*slot*]*ve_port options arguments*

portcfg action [*slot*][*ge_port options arguments*

Description

Use this command to manage port configuration parameters on FC ports as well as on Gigabit Ethernet (GbE) ports on the Brocade 7840 switch and on the Brocade FX8-24 and Brocade SX6 blades.

You must use this command in a manner that honors the platform-specific differences in command syntax and behavior. Some command options are not available on all platforms. Others behave differently depending on the platform on which they are executed. Use the following section headings to navigate this page.

Commands supported on all platforms

- **portcfg mirrorport** - Configure a mirror port on the local FC port.
- **portcfg rscnsupr** - Manage registered state change notification (RSCN) suppression on the local port.

Configure IP interfaces on the Brocade 7840 switch, Brocade FX8-24 and Brocade SX6 blades.

- **portcfg ipif** - Configure the local IP interfaces.
- **portcfg iproute** - Configure a static route on the IP interface.
- **portcfg vlantag** - Manage the IP interface VLAN configuration for FCIP. This command is supported on Brocade FX8-24 blade.
- **portcfg ipsec-policy** - Include the IPsec policy. Currently supported on the Brocade 7840 switch and Brocade SX6 blade only.
- **portcfg lag** - Configure a link aggregation group (LAG) for a GE LAN port. Supported on the Brocade 7840 switch and Brocade SX6 blade only.
- **portcfg tcl** - Configure the Traffic Control List (TCL). Supported on the Brocade 7840 switch and Brocade SX6 blade only.
- **portcfg mgmtif** - Creates the inband management interfaces.
- **portcfg mgmroute** - Creates the routes for the inband management interfaces.
- **portcfg autoneg** - Configure autonegotiation settings for 1 GbE ports.
- **portcfg sla** - Configure an Service Level Agreement (SLA) session. Supported on the Brocade 7840 switch and Brocade SX6 blade only.

Configure tunnels, circuits, and management interfaces on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades.

- **portcfg fcipunnel** - Create, modify, and delete Fibre Channel over IP (FCIP) tunnels
- **portcfg fcipcircuit** - Create, modify, and delete FCIP circuits.

To display the command usage on the switch, execute **portcfg action** without any further arguments.

Notes IPv6 addresses are supported except for inband management. IP Security enabled FCIP Tunnels cannot use IPv6 circuits on FX8-24 blade.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Some of the features supported by this command may require a license.

The Fabric OS port configuration commands are not supported on FCoE ports.

Function Commands supported on all platforms

Synopsis `portcfg action [slot/]port arguments`

Description Use this command to configure the following parameters on a local FC port.

- **portcfg mirrorport** - Configure a mirror port on the local port.
- **portcfg rscnsupr** - Manage registered state change notification (RSCN) suppression on the local port.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

port[-port]

Specifies a single port or a range of ports, for example, 3-22 or 1/3-8. port ranges are supported only with the **rscnsupr** option.

rscnsupr

Manages Registered State Change Notification (RSCN) suppression on the local port. RSCN suppression is configurable only on FC ports. The syntax for **portCfg rscnsupr** is as follows:

portcfg rscnsupr [slot/]port[-port] mode

The following modes are supported with the **rscnsupr** option:

--disable

Disables the configuration. When disabled, device changes on the port generate an RSCN to all other end devices that are zoned with this one. By default, RSCN suppression is disabled on all ports.

--enable

Enables the configuration. When enabled, any device changes on the port will not generate an RSCN to any other end device.

mirrorport

Configures a mirror port on the local port. The port mirroring feature reroutes data frames between two devices to the mirror port. Port mirroring can aid in troubleshooting common FC end-to-end communication problems. The command prompts for confirmation that the specified port be enabled as a mirror port. Once a port is configured as a mirror port, the port can only be used as part of a mirror connection. Port Mirroring is not supported over FCIP links (VE tunnels). The port mirroring feature is not supported on the Gen6 platforms.

The syntax for **portCfg mirrorport** is as follows:

portcfg mirrorport [*slot/*]*port mode*

Valid modes for **mirrorport** include the following:

--disable

Disables the configuration. When disabled, a port cannot be a mirror port.

--enable

Enables the configuration. When the mirror port feature is enabled to a port, a mirror connection can use this port to mirror traffic.

Examples To enable a mirror port configuration:

```
switch:admin> portcfg mirrorport 2/4 --enable
Please confirm enable of Mirror Port (Y,y,N,n):[n] y
```

To configure a range of ports as RSCN-suppressed:

```
switch:admin> portcfg rscnsupr 2/4-7 --enable
```

Function Configure extension-related parameters on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades.

Synopsis **portcfg** *action* [*slot/*]*ge_port* | *name arguments*

Description Use this command to configure the local IP interfaces and static routes on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades. You must configure the local IP interfaces before you can create and configure FCIP tunnels. You can also create a VLAN configuration at the IP interface on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades.

- **portcfg ipif** - Configure the local IP interfaces.
- **portcfg iproute** - Configure a static route on the IP interface.
- **portcfg vlantag** - Manage the IP interface VLAN configuration for FCIP. This command is supported on Brocade FX8-24 blade.
- **portcfg ipsec-policy** - Include the IPsec policy. Currently supported on the Brocade 7840 switch and Brocade SX6 blade only.
- **portcfg lag** - Configure a link aggregation group (LAG) for a GE LAN port. Supported on the Brocade 7840 switch and Brocade SX6 blade only.
- **portcfg tcl** - Configure the Traffic Control List (TCL). Supported on the Brocade 7840 switch and Brocade SX6 blade only.
- **portcfg mgmtif** - Creates the inband management interfaces.
- **portcfg mgmtroute** - Creates the routes for the inband management interfaces.
- **portcfg autoneg** - Configure autonegotiation settings for 1 GbE ports.
- **portcfg sla** - Configure an Service Level Agreement (SLA) session. Supported on the Brocade 7840 switch and Brocade SX6 blade only.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

ge_port

Specifies the number of the GbE port to be configured. The GbE ports are numbered ge0 - ge9 on the Brocade FX8-24 blade and ge0 - ge17 on Brocade 7840 switch and Brocade SX6 extension blade.

The two 10GbE ports on the Brocade FX8-24 blade are numbered xge0 and xge1.

The Brocade 7840 switch and the Brocade SX6 extension blade has two 40GbE ports labeled ge0 and ge1, and 16 1GbE or 10GbE ports labeled ge2-ge17. For Brocade 7840 switch and the Brocade SX6 extension blade, specify GbE port number along with the DP number, for example, **portcfg ipif ge0.dp0**. The valid DP numbers are dp0 and dp1.

For Brocade 7840 and the Brocade SX6 extension blade, specify the switch virtual interface (SVI) LAN port in "lan.dp#" format, for example, **portcfg ipif lan.dp0**.

Use the **switchShow** command for a list of valid ports.

ipif

Defines the IP interface for both ports of a tunnel.

The IP network connection is configured by defining IP interfaces for origin and destination virtual ports, and then defining one or more IP routes to connect them. The syntax for **portCfg ipif** is as follows:

```
portcfg ipif [slot/]ge_port option args [optional_args]
```

Valid options and arguments for **ipif** include the following:

```
create src_ipaddr netmask mtu [mask] mtu_size vlan [vlan_id]
```

Creates IP interfaces. Specify the following:

src_ipaddr

Specifies source IP address in either IPv6 or IPv4 format:

src_IPv6_addr/prefix_len

Specifies the source IPv6 address of the port if IPv6 is used. The address must be an IPv6 global, unicast address, followed by a prefix. This is used for IPv6 addresses instead of a netmask. The *prefix_len* operator is required. Refer to the *Fabric OS Extension Administrator's Guide* for more information on IPv6 rules and restrictions.

src_IPv4_addr netmask [mask]

Specifies the source IPv4 address of the port, if IPv4 is used. If an IPv4 address is used, the subnet mask must be specified as well (in a.b.c.d. format). The CIDR notation for IPv4 addresses can also be used.

mtu [mtu_size]

Specifies the maximum transmission unit size. The permitted range is 1260 Bytes to 1500 Bytes for Brocade 7800 switch and FX8-24 blade, and 1280 Bytes to 9216 Bytes or 'auto' to enable PMTU discovery for Brocade 7840 switch and the Brocade SX6 extension blade. This operand is optional.

vlan [vlan_id]

Specifies the VLAN ID. This operand is optional and currently supported on the Brocade 7840 and the Brocade SX6 extension blade only.

-x | --crossport

Configures an interface for a failover crossport. A crossport is defined as the nonlocal (secondary) XGE port to a VE_Port group that provides failover capabilities for the tunnel configured on the local XGE port. For VE_Ports 12-21, xge1 is the local XGE port and xge0 is the crossport. For VE_Ports 22-31, xge0 is the local XGE port and xge1 is the crossport. This operand is optional and supported only on the Brocade FX8-24 XGE ports.

delete ipaddr

Deletes IP interfaces. Specify the IPv6 or IPv4 address of the virtual port.

iproute

Defines static IP routes on a GbE port or on a 10GbE port. After defining the IP interface on the remote switch, you can define destination routes for an interface. You cannot use this command to configure routes for inband management; use **portCfg mgmtroute** instead.

The syntax for **portcfg iproute** is as follows:

```
portcfg iproute [slot/]ge_port option args [optional_args]
```

Valid options and arguments for **iproute** include the following:

create dest_ipaddr netmask [mask] [gateway_router]

Creates an IP route. Specify the following:

dest_ipaddr

Specifies the destination IP address in either IPv6 or IPv4 format:

dest_IPv6_addr/prefix_len

Specifies the destination IPv6 address of the port, if IPv6 is used. The address must be an IPv6 global, unicast address, followed by a prefix. This is used for IPv6 addresses instead of a netmask. The *prefix_len* operand is required.

dest_IPv4_addr netmask [mask]

Specifies the destination IPv4 address of the port, if IPv4 is used. If an IPv4 address is used, the subnet mask must be specified as well. Use a.b.c.d. format. The CIDR notation for IPv4 addresses can also be used.

gateway_router

Specifies the IP address of an IP router that can route packets to the destination virtual port IP address. The gateway address must be on the same IP subnet as one of the port IP addresses. This operand is optional with IPv6 addresses. If not specified, the *gateway_router* learned from the Neighbor Discovery protocol is used.

-x | --crossport

Configures an IP route for a failover crossport. A crossport is defined as the nonlocal (secondary) XGE port to a VE_Port group that provides failover capabilities for the tunnel configured on the local XGE port. For VE_Ports 12-21, xge1 is the local XGE port and xge0 is the crossport. For VE_Ports 22-31, xge0 is the local XGE port and xge1 is the crossport. This operand is optional and supported only on the Brocade FX8-24 XGE ports.

delete ipaddr

Deletes IP routes for a specified IPv4 or IPv6 address. Valid *ipaddr* values include the following:

ipaddr

Specifies IP address in either IPv6 or IPv4 format:

IPv6_addr/prefix_len

Specifies the IPv6 address of the virtual port, if IPv6 is used. The address must be an IPv6 global, unicast address, followed by a prefix. The *prefix_len* operand is required.

dest_IPv4_addr netmask

Specifies the destination IPv4 address of the virtual port, if IPv4 is used. If an IPv4 address is used, the subnet mask must be specified as well. Use a.b.c.d. format or the CIDR notation for IPv4 addresses.

vlantag

Use this command to maintain the VLAN tag table associated with a particular network interface. This command is supported on Brocade FX8-24 blade. This table is mainly used during ingress processing to filter VLAN tagged frames. You can configure multiple VLAN IDs per IP interface; however, if you provide a destination address, there cannot be two entries to the same destination, including 0.0.0.0. Egress frames already marked as VLAN tagged (FCIP circuit-level tagging) take precedence over entries in this table.

This command supports **add** and **delete** options only. To modify a table entry, it must first be deleted, then added with different configuration parameters. The syntax for **portcfg vlantag** is as follows:

portcfg vlantag [*slot*]/*ge_port mode arguments*

Valid modes for **vlantag** include the following:

add *ipif_addr vlan_id l2cos* [*dst_ipaddr*][*optional_args*]

Adds an entry to the VLAN tag table.

delete *ipif_addr vlan_id* [*dst_ipaddr*]

Deletes an entry from the VLAN tag table. Valid arguments for **add** and **delete** include the following:

ipif_addr

Specifies the locally defined interface address in IPv6 or IPv4 format.

vlan_id

Specifies the VLAN ID used for this tag. The range is 1 to 4094.

l2cos

Specifies Layer 2 Class of Service/Priority, as defined in the IEEE 802.1p specification. The range is 0 to 7. Valid only with the **add** option.

dst_ipaddr

Specifies an optional destination IP address (IPv4 or IPv6). All packets destined for this IP address are tagged accordingly. If a destination IP address is not specified, all packets not already tagged will be tagged. The default is 0.0.0.0.

optional_args

Optional VLAN tagging parameters include the following:

-x | --crossport

>Applies the VLAN tag to a crossport interface. Specifying this parameter allows the VE_Ports 12-21 to use the IP interface with this vlantag. This operand is optional and valid only and supported only on the Brocade FX8-24 XGE ports.

-m | --mgmt

Specifying this parameter applies the VLAN tag to the inband management interface.

ipsec-policy

Creates an IPsec policy. This option is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.

This command supports **create**, **modify**, and **delete** options only. The syntax for **portcfg ipsec-policy** is as follows:

portcfg ipsec-policy name option [args]

Valid options and arguments for **ipsec-policy** include the following:

create name [-k | --preshared-key key]

Creates the IPsec policy.

name

Specifies the name for the IPsec policy.

-k | --preshared-key key

Specifies the preshared key to be used for authentication. This operand is optional.

modify name [-k | --preshared-key key]

Modifies the IPsec policy.

name

Specifies the name for the IPsec policy.

-k | --preshared-key key

Specifies the preshared key to be used for authentication. This operand is optional.

delete name

Deletes the IPsec policy.

name

Specifies the name for the IPsec policy.

lag name

Configures a link aggregation group (LAG) for a GE LAN port, modifies a LAG, and displays LAG status. The syntax for **portcfg lag** is as follows:

portcfg lag lag_name [arguments]

Valid options and arguments for **lag** include the following:

--create [ge_port]

Creates a LAG. You can specify a single port or a range of ports, for example, ge2 or ge2-ge4 or ge2,ge4. Use the **switchShow** command for a list of valid ports.

--add [ge_port]

Adds the specified GE ports to the LAG. You can specify a single GE port or a range of GE ports to be added to the LAG, for example, ge2 or ge2-ge4 or ge2,ge4.

--remove [ge_port]

Removes the specified GE ports from the LAG. You can specify a single port or a range of ports to be removed from the LAG, for example, ge2 or ge2-ge4 or ge2,ge4.

--delete

Deletes the specified LAG.

--set -speed speed

Configures port speed for the LAG port members. The valid speeds are 1G and 10G.

--enable -autoneg
Enables the auto negotiation mode for the LAG port members.

--disable -autoneg
Disables the auto negotiation mode for the LAG port members.

tcl

Creates a TCL.

This command supports **create**, **modify**, and **delete** options only. The syntax for **portcfg tcl** is as follows:

portcfg tcl name option [args]

Valid options and arguments for **tcl** include the following:

name

Specifies the name of the TCL. The TCL name can be up to 31 characters long and cannot contain special characters.

create name [args]
Creates a TCL.

modify name [args]
Modifies a TCL.

-p | --priority value
Sets the priority ID for the TCL. The range is from 1 through 65534.

--admin-status enable | disable
Enables or disables the TCL.

--action allow | deny | [slot/]dp#-deny
Specifies the permit or deny action associated with the TCL.

-t | --target VE[-pri]
Sets the target virtual interface and priority for the TCL.

-S | --src-addr ipaddr[/prefix_len]
Sets the source IP address input filter for the specified TCL.

--src-mask ipaddr
Sets the source IP address mask input filter for the specified TCL.

-D | --dst-addr ipaddr[/prefix_len]
Sets the destination IP address input filter for the specified TCL.

--dst-mask ipaddr
Sets the destination IP address mask input filter for the specified TCL.

--proto-port ipaddr
Sets the protocol port input filter for the specified TCL.

--proto-app app_name
Sets the application input filter for the specified TCL.

--dscp value
Set the DSCP input filter for the specified TCL. The range is from 0 through 63.

- v | --vlan *value***
Set the VLAN input filter for the specified TCL. The range is from 0 to 4095.
- l2cos *value***
Sets the Layer 2 CoS input filter for the specified TCL.
- l4proto *num | name***
Sets the Layer 4 protocol input filter for the specified TCL. Valid values for *name* are TCP, UDP, ICMP, VRRP and the valid range for *num* is from 0 through 255, or 'none' to clear the Layer 4 protocol input filter.
- rst-propagation enable | disable**
Enables or disables end-to-end RST propagation for the specified TCL.
- segment-preservation enable | disable**
Enables or disables segment preservation for the specified TCL.
- non-terminated enable | disable**
Enables or disables non-terminate traffic handling for the specified TCL.

mgmtif

Configures an inband management interface. A maximum of one interface is supported per GbE Port. Up to a total of 11 interfaces for the Brocade FX8-24 (depending on blade operating mode). The interfaces must each exist on a unique network, separate from the other inband management interfaces and from all other interfaces on the CP. The syntax for **mgmtif** is as follows:

mgmtif *ge_port [options] arguments*

Valid options and arguments for **mgmtif** include the following:

create

Creates a specified IP Address for the inband management interface. The IP address must be unique.

IPv4_address

Specifies an IP address for the inband management interface in IPv4 format, followed by the subnet mask.

netmask

Specifies the subnet mask for the IPv4 address in a.b.c.d format.

MTU

Specifies the MTU for the inband management interface. The valid range is 1260 to 1500. This operand is optional.

delete

Deletes the specified IP Address for the inband management interface.

IPv4_address

Specifies the IP address to be deleted.

disable

Disables the inband management interface at the specified GbE Port.

| | |
|------------------|--|
| enable | Re-enables the inband management interface at the specified GbE Port after it has been disabled. The interface is by default enabled when it is created. |
| mgmtroute | <p>Configures the management route for the inband management interface (needed only if the management interfaces are on different subnets). The syntax for mgmtroute is as follows:</p> <p>mgmtroute ge_port [options] destination_ip_address</p> <p>Valid options and arguments for mgmtroute include the following:</p> |
| create | <p>Creates a management route for a specified destination IP Address. The following operands are required:</p> <p><i>dest_IPv4_addr</i> Specifies the destination IP address for the management route in IPv4 format.</p> <p><i>netmask</i> Specifies the subnet mask for the IPv4 address in a.b.c.d format.</p> <p><i>gateway</i> Specifies the IP address of an IP router that can route packets to the destination IP address.</p> |
| delete | <p>Deletes a management route for a specified destination IP Address. The following operands are required.</p> <p><i>dest_IPv4_addr</i> Specifies the destination IP address for the management route in IPv4 format.</p> <p><i>netmask</i> Specifies the subnet mask for the IPv4 address in a.b.c.d format.</p> |
| sla | <p>Creates, modifies, or deletes an SLA session. For more information on the SLA feature, refer to the <i>Fabric OS Extension Administration Guide</i>.</p> <p>This command supports create, modify, and delete options only. The syntax for portcfg sla is as follows:</p> <p>portcfg sla name option [args]</p> <p>Valid options and arguments for sla include the following:</p> |
| create | Creates an SLA session with the specified name. You must create an SLA session at each end of the circuit, but the session names need not match. |
| modify | Modifies the specified SLA session. |
| <i>name</i> | Specifies the name of the SLA. |

--loss percentage
Sets the packet-loss percentage. The valid range is from 0.05 through 5.0 percentage.

--runtime min
Sets duration time for the test to run. The valid range is from 1 through 1440 minutes. The default value is 5 minutes.

--timeout min | none
Sets the timeout duration for the test. If the timeout value is reached during the SLA session, the session is terminated and the circuit is put into service. Specify **none** to run the test until the runtime and packet-loss values are met. The valid range is from 1 through 2880 minutes.

delete name
Deletes the specified SLA session.

Examples To create an IP interface using IPv4:

```
switch:admin> portcfg ipif ge0 create \
    192.169.0.20 netmask 255.0.0.0 mtu 1500
Operation Succeeded
```

To create an SVI LAN port on a DP:

```
switch:admin> portcfg ipif lan.dp0 create 10.0.0.1/24 vlan 100
Operation Succeeded.
switch:admin> portcfg ipif lan.dp0 create 10.0.1.1/24 vlan 200
Operation Succeeded.
```

```
switch:admin> portshow ipif
```

| Port | IP Address | / Pfx | MTU | VLAN | Flags |
|----------|----------------|-------|------|------|-------|
| ge4.dp0 | 192.168.60.20 | / 24 | 1500 | 0 | U R M |
| ge17.dp0 | 192.168.10.107 | / 24 | 1500 | 0 | U R M |
| lan.dp0 | 10.0.0.1 | / 24 | 1500 | 100 | U R M |
| lan.dp0 | 10.0.1.1 | / 24 | 1500 | 200 | U R M |

Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

To create an IP interface using IPv6 with a prefix:

```
switch:admin> portcfg ipif ge0 create 2000::22/64 1500
Operation Succeeded
```

```
switch:admin> portshow ipif ge0
```

| Port | IP Address | / Pfx | MTU | VLAN | Flags |
|------|--------------|-------|------|------|-------|
| ge0 | 192.180.0.20 | / 24 | 1500 | n/a | U R M |
| ge0 | 192.168.0.21 | / 24 | 1500 | n/a | U R M |
| ge0 | 2000::10 | / 64 | 1500 | n/a | U R M |

Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

To create an interface for a failover crossport on a Brocade FX8-24:

1. Configure the interface for the local XGE port (xge0).

```
switch:admin> portcfg ipif 8/xge0 create \
    192.169.0.20 netmask 255.255.255.0 mtu 1500
Operation Succeeded
```

2. Configure the interface for the nonlocal XGE port (xge1).

```
switch:admin> portcfg ipif 8/xge1 create \
    192.169.0.20 netmask 255.255.255.0 -x
Operation Succeeded
```

3. The IP address 192.168.11.20 will be available for circuits on VE_Ports 22-31 on the Brocade FX8-24 extension blade. The output from **portshow ipif** for the same address shows the crossport tag.

```
switch:admin> portshow ipif 8/xge0
```

| Port | IP Address | / Pfx | MTU | VLAN | Flags |
|--------|---------------|-------|------|------|---------|
| 8/xge0 | 192.168.10.20 | / 24 | 1500 | n/a | U R M |
| 8/xge0 | 192.168.11.21 | / 24 | 1500 | n/a | U R M X |

Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

To delete an IP interface:

```
switch:admin> portcfg ipif ge0 delete 192.168.10.20
Operation Succeeded
```

To delete an IP interface for a failover crossport:

```
switch:admin> portcfg ipif 8/xge0 delete 192.168.11.20 -x
Operation Succeeded
```

To create a static IP route using an IPv4 destination address, a netmask, and a gateway address:

```
switch:admin> portcfg iproute ge0 create \
    192.42.0.0 netmask 255.255.255.0 192.168.0.250
Operation Succeeded
```

To create a static IP route using IPv6:

```
switch:admin> portcfg iproute ge0 create \
    2010::/64 2000::1:250
Operation Succeeded
```

```
switch:admin> portshow iproute ge0
```

| Port | IP Address | / Pfx | Gateway | Flags |
|------|-------------|-------|---------|-------|
| ge0 | 192.168.0.0 | / 24 | * | U C |
| ge0 | 2000:: | / 64 | * | U C |
| ge0 | 2000:::10 | / 128 | * | U H L |

```

ge0          2000::1:250          / 128 *          U H L
ge0          2010::              / 64  2000::1:250  U G S
-----

```

```

Flags: U=Usable G=Gateway H=Host C=Created(Interface)
       S=Static L=LinkLayer X=Crossport

```

To create an IP route for a failover crossport on a Brocade FX8-24:

1. Configure the iproute for the local XGE port (xge0).

```

switch:admin> portcfg iproute 8/xge0 create 2.2.2.0 \
               netmask 255.255.255.0 192.168.11.250
Operation Succeeded

```

2. Configure the interface for the non-local XGE port (xge1).

```

switch:admin> portcfg ipif 8/xge1 create 1.1.1.0 \
               netmask 255.255.255.0 192.168.11.250 -x
Operation Succeeded

```

3. The IP route will be available for circuits on VE ports 22-31. The output from **portshow iproute** for the same route shows the crossport tag.

```

switch:admin> portshow iproute 8/xge0

```

| Port | IP Address | / Pfx | Gateway | Flags |
|--------|---------------|-------|---------|---------|
| 8/xge0 | 192.168.0.0 | / 24 | * | U C |
| 8/xge0 | 192.168.0.111 | / 32 | * | U H L |
| 8/xge0 | 192.168.1.0 | / 24 | * | U C X |
| 8/xge0 | 192.168.1.111 | / 32 | * | U H L X |

```

-----
Flags: U=Usable G=Gateway H=Host C=Created(Interface)
       S=Static L=LinkLayer X=Crossport

```

To delete the crossport IP route:

```

switch:admin> portcfg iproute 8/xge1 delete \
               1.1.1.0 netmask 255.255.255.0 -x
Operation Succeeded

```

To delete a static IP route using an IPv4 address:

```

switch:admin> portcfg iproute ge0 delete \
               172.16.0.0 netmask 255.255.0.0
Operation Succeeded

```

To create a network-wide permanent VLAN tag entry with a VLAN ID of 200 and an L2 CoS value of 5 (no destination address):

```

switch:admin> portcfg vlantag ge2 add \
               192.168.2.10 200 5
Operation Succeeded

```

To display the VLAN tag configuration (for an explanation of the flags, refer to **portShow**):

```
switch:admin> portshow vlantag ge2

Port: ge2
Interface Address          Destination Address      VlanId L2CoS  Flags
-----
192.168.2.10              0.0.0.0                 200    5      Perm Net
192.168.0.20              192.168.0.10           100    0      Perm
192.168.0.21              192.168.0.11           200    0      Perm
2000::20                  2000::10                300    0      Perm
2000::21                  2000::11                400    0      Perm
```

To create a VLAN tag entry for a failover crossport interface with a VLAN ID of 200 and an L2 CoS value of 1 and to display the configuration:

```
switch:admin> portcfg vlantag 8/xge0 add 192.168.11.20 200 1 -x
Operation Succeeded
switch:admin> portshow vlantag 8/xge0
Port: 8/xge0
Interface Address          Destination Address      VlanId L2CoS  Flags
-----
192.168.11.20             0.0.0.0                 200    1      Perm Net (crossport)
```

To delete a VLAN tag entry:

```
switch:admin> portcfg vlantag ge2 delete 192.168.2.10 200
Operation Succeeded
```

To delete a VLAN tag entry for a crossport:

```
switch:admin> portcfg vlantag 8/xge0 delete 192.168.11.20 200 -x
Operation Succeeded
```

To create an IPsec policy for the Brocade 7840 switch:

```
switch:admin> portcfg ipsec-policy myPolicy1 create -k "some test key"
Operation Succeeded
```

To view the IPsec policy status for the Brocade 7840 switch and the Brocade SX6 extension blade:

```
switch:admin> portshow ipsec-policy all --ike
IPsec Policy  OpStatus IKECnt  Pre-Shared Key
  IKE-ID      OpStatus RefCnt  Local IP Address  Remote IP Address
-----
myPolicy1    NotInUse   0    some test key
```


To create a LAG:

```
switch:admin> portcfg lag lag0 --create
Operation Succeeded
```

To add ports to a LAG and verify the configuration:

```
switch:admin> portcfg lag lag0 --add ge2-ge4,ge7
Operation Succeeded
switch:admin> portshow lag
LAG Name   State      Port Count  Member Ports
-----
lag0       Online     4           GE2,GE3,GE4,GE7
-----
```

To remove a port from a LAG:

```
switch:admin> portcfg lag lag0 --remove ge7
Operation Succeeded
```

To create a TCL:

```
switch:admin> portcfg tcl hostAtoB create --admin enable \
--target 24 --src-addr 10.0.0.0/8 --priority 10
Operation Succeeded
```

To delete an inband management interface:

```
switch:admin> portcfg mgmtif ge0
delete 192.168.3.10 255.255.255.0
```

To delete a management route:

```
switch:admin> portcfg mgmtroute ge0 delete 192.168.3.0 255.255.255.0 192.168.2.250
```

To disable and re-enable an inband management interface:

```
switch:admin> portcfg mgmtif ge0 disable
switch:admin> portcfg mgmtif ge0 enable.
```

To configure a VLAN tag for an inband management port:

```
switch:admin> portcfg vlantag 10/xge0 add 172.16.30.40 30 0 --mgmt
Operation Succeeded
```

To configure VLAN tags for all ports on a switch:

```
switch:admin> portshow vlantag all
Port: 10/xge0
Interface Address          Destination Address      VlanId  L2CoS  Flags
-----
172.16.30.80                172.16.50.80            30      0      Perm
172.16.30.40                0.0.0.0                  30      0      Perm Net Mgmt
-----
```

To enable the autonegotiation of a 1GbE port:

```
switch:admin> portcfg autoneg 1/ge2 --enable
Operation Succeeded
```

```
switch:admin> portshow autoneg 1/ge2
Auto-Negotiation Configuration:
Port: 1/ge2
Auto-Negotiate: Enabled
```

To create an SLA session:

```
switch:admin> portcfg sla netA create --loss 0.5 \
--runtime 15 --timeout 60
Operation Succeeded
```

Function Configure tunnels, circuits, and management interfaces on the Brocade 7840 switch and Brocade FX8-24 and Brocade SX6 extension blades.

Synopsis `portcfg action [slot]ve_port options arguments`

Description Use this command to configure FCIP tunnels and FCIP circuits on the Brocade 7840 switch, and the Brocade FX8-24 and Brocade SX6 extension blades only. The FCIP tunnels on the local and remote ports act as Virtual E_Ports (VE_Ports) connecting the local and remote fabrics. The following operations can be performed with this command:

- **portcfg fciptunnel** - Create, modify, and delete Fibre Channel over IP (FCIP) tunnels.
 - Configure a tunnel with the FCIP FastWrite feature.
 - Enable or disable Tape Pipelining.
 - Configure the compression options.
 - Configure an IPsec-enabled tunnel.
 - Configure VLAN tagging on the default FCIP circuit. This option is supported on Brocade FX8-24 blade only.
 - Configure Differentiated Services Code Point (DSCP) markings.
 - Set the committed rate or minimum and maximum rates for the default circuit.
 - Configure a tunnel for FICON emulation.
- **portcfg fcipcircuit** - Create, modify, and delete additional FCIP circuits.
 - Set or modify the committed rate or minimum and maximum rates for the circuit.
 - Configure or change VLAN tagging on the default circuit.
 - Set or modify additional circuit parameters.

Notes You can configure up to 20 tunnels on the Brocade 7840 switch, and on the Brocade FX8-24 and Brocade SX6 extension blades.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

ve_port

Specifies the number of the VE_Port associated with the tunnel. The VE_Ports are numbered 24-43 on the Brocade 7840, 16-35 on the Brocade SX6, and 12-31 on the Brocade FX8-24.

fciptunnel

Creates, modifies, or deletes an extension tunnel.

Note the following port mapping rules:

| | 7840 | SX6 | FX8-24 |
|-----------|-------|-------|-----------------------|
| DP0 ports | 24-33 | 16-25 | 22-31 (xge0) |
| DP1 ports | 34-43 | 26-35 | 12-21 (ge0-ge9, xge1) |

The valid VE ports on the Brocade 7840 switch are VE_Ports 24-43. In 10VE mode, only VE_Ports 24-28 and VE_Ports 34-38 are supported. In 20VE mode all VE ports are supported.

The syntax for **portcfg fciptunnel** is as follows:

```
portcfg fciptunnel [slot/]ve_port options [tunnel_arguments] [circuit_arguments]
```

The following *options* are supported with **fciptunnel**:

create [tunnel_arguments][circuit_arguments]

Creates an extension tunnel and, optionally a single default circuit. If no circuit arguments are specified, no extension circuit will be created. While it is possible to create a tunnel without a default circuit, it is an unlikely scenario (for example, for configuration staging purposes). In most cases, you will create a tunnel with at least one configured circuit. In this case, you must specify a remote and local IP address for the circuit as well as a committed rate (or alternately, a minimum and maximum committed rate) to configure the default circuit. The default circuit created with the tunnel is automatically assigned the circuit ID 0. You can modify the default circuit with the **fciptunnel modify** command. To add additional circuits to an existing tunnel, use the **fcipcircuit create** command.

modify [tunnel_arguments][circuit_arguments]

Modifies the properties of an existing extension tunnel. To modify a tunnel, you must specify at least one of the tunnel or circuit parameters for the command to be effective. Any circuit attribute you change with the **fciptunnel modify** command affects only the circuit 0. All other circuits remain unchanged. To modify a circuit other than circuit 0, you must use the **fcipcircuit modify** command.

In order to create a default circuit 0 with the tunnel, you must at least specify the following circuit arguments with **fciptunnel create**:

```
-D | --remote-ip-address remote_ip_address
```

Specifies the IP address for the remote end of the extension circuit.

```
-S | --local-ip-address local_ip_address
```

Specifies the IP address for the local end of the extension circuit.

-b | --min-comm-rate *value* -B | --max-comm-rate *value*

You may set a minimum and a maximum for the committed rate to configure the tunnel for Adaptive Rate Limiting (ARL), which allows for a more effective sharing of bandwidth between applications. For Brocade FX8-24, the valid range is 10,000 Kbps to 1,000,000 Kbps for 1G ports and 10,000 Kbps to 10,000,000 Kbps for 10G ports. For Brocade 7840, the valid range is 20,000 Kbps to 10,000,000 Kbps. The maximum committed rate can be no larger than five times the minimum committed rate.

In the previous releases, the maximum and minimum rates had to be configured to the same value for circuits using the XGE ports on the FX8-24 blade. You can configure different values for minimum and maximum rates for each circuit of a tunnel using the XGE ports.

Optional *tunnel_arguments* for **fciptunnel create** and **modify** include the following. Tunnel parameters are by default disabled. To change the default (for example, enabling FastWrite) with **fciptunnel create**, specify the parameter only. To modify any of these parameters with **fciptunnel modify**, specify the parameter and one of the values in square brackets.

-f | --fastwrite [disable | enable]

Disables or enables the FCIP FastWrite on the specified extension tunnel.

-t | --tape-pipelining [disable | enable | write-only]

Configures Open Systems Tape Pipelining on the specified extension tunnel. By default, OSTP is disabled .

When using this operand with **fciptunnel create**, specify one or more of the following operands:

-t | --tape-pipelining

Enables write-read Tape Pipelining (FCIP FastWrite must also be enabled.)

-N | --no-read-pipelining

Disables tape read-pipelining. This operand is valid only with **fciptunnel create** and **-t** must be specified enabling the feature. The combination of **-t** and **-N** effectively enables tape write-pipelining. This operand is supported on the Brocade FX8-24 blade only.

When using this operand with **fciptunnel modify**, specify one or more of the following:

-t | --tape-pipelining *mode*

Modifies the Open System Tape Pipelining configuration. Specify one of the following modes:

disable

Disables Tape Pipelining

enable

Enables write-read Tape Pipelining (FCIP FastWrite must also be enabled).

write-only

Enables write-only Tape Pipelining (FCIP FastWrite must also be enabled).

- c | --compression *compression_level***
 Configures compression on the specified FCIP tunnel. By default, compression is disabled. Specify one of the following values on the FX8-24 blades:
- none**
Compression disabled
 - hardware**
Enable standard compression.
 - moderate**
Enable moderate compression.
 - aggressive**
Enable aggressive compression.
 - auto**
enable auto compression. Automatically adjusts compression level based on the maximum configured tunnel bandwidth. To enable this feature upgrade both ends of the tunnel to Fabric OS v7.0.0. Based on total effective tunnel bandwidth, the compression level will be adjusted as follows:
 - Aggressive**
Bandwidth less than 512 Mbps
 - Moderate**
Bandwidth more than 512 Mbps and less than 2 Gbps.
 - Standard**
Bandwidth more than 2Gbps.
- Specify one of the following values on the Brocade 7840 switch and the Brocade SX6 extension blade:
- none**
Compression disabled.
 - fast-deflate**
Enable fast-deflate compression.
 - deflate**
Enable deflate compression.
 - aggr-deflate**
Enable aggressive deflate compression level.
- fc-compression *compression_level***
 Configures FC compression on the specified FCIP tunnel. By default, compression is disabled. This operand is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only. Specify one of the following values for *compression_level*:
- aggr-deflate**
Enable aggressive deflate compression level.

- default**
Enable default compression level.
- deflate**
Enable deflate compression.
- fast-deflate**
Enable fast-deflate compression.
- none**
Compression disabled.
- ip-compression *compression_level***
Configures IP compression on the specified FCIP tunnel. By default, compression is disabled. This operand is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only. Specify one of the following values for *compression_level*:
- aggr-deflate**
Enable aggressive deflate compression level.
- default**
Enable default compression level.
- deflate**
Enable deflate compression.
- none**
Compression disabled.
- ipext enable | disable**
Enables or disables IP Extension capability to a tunnel. This parameter is supported only on Brocade 7840 switch and the Brocade SX6 extension blade.
- T | --tperf [0|1]**
Disables (0) or enables (1) TPerf test mode. Refer to **portCmd** help for more information regarding TPerf monitoring. This operand is supported on the Brocade FX8-24 blade only.
- n | --remote-wwn *remote-wwn***
Specifies the WWN of the remote FC entity.
- d | --description *string***
Specifies a description for the specified tunnel.
- i | --ipsec [disable | enable]**
Disables or enables Internet Protocol Security (IPsec) on the specified tunnel. Circuits that fall underneath a tunnel inherit the IPsec tunnel attributes. IPsec uses a predefined policy with IKEv2 for key negotiation, ESP transport mode for IPsec, and AES with 256-bit keys for Encryption (AES-GCM-ESP). You must specify a key with this option. On the Brocade FX8-24, IPsec is supported only on VE_Ports 12-21. This operand is supported on the Brocade FX8-24 blade only.
- K | --key *presared_key***
Specifies the preshared key to be used for authentication. Specify a string of alphanumeric characters 32 bytes in length. This argument must be used together with **--ipsec**. This argument is supported on the Brocade FX8-24 blade only.

- i | --ipsec [*policy_name* | none]**
Sets the Internet Protocol Security (IPsec) on the specified tunnel to use the specified IPsec Policy, or to disable IPsec for the tunnel if the 'none' operand is used. This argument is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- p | --distribution *mode[,ratio,...]***
Sets tunnel bandwidth distribution mode. The valid modes are protocol, priority, and default. Optionally, you can change the default bandwidth allocation ratio for Fibre Channel (FC) and IP Extension (IP) traffic in the tunnel. This operand is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- Q | --fc-qos-ratio *high,med,low***
Sets QoS percentages for FC priorities. The default values are 50% (high), 30% (medium), and 20% (low). Each priority can have a minimum of 10%. The sum of percentages must equal 100%. This operand is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- I | --ip-qos-ratio *high,med,low***
Sets QoS percentages for IP priorities. The default values are 50% (high), 30% (medium), and 20% (low). Each priority can have a minimum of 10%. The sum of percentages must equal 100%. This operand is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- q | --qos-bw-ratio *high,med,low* | default**
Sets QoS percentages for all QoS priorities. The default values are 50% (high), 30% (medium), and 20% (low). Each priority can have a minimum of 10%. The sum of percentages must equal 100%. Both ends of the tunnel must have the same QoS priority settings.
- F | --ficon [*disable* | *enable*]**
Enables or disables FICON emulation on the specified extension tunnel. Optional FICON arguments for **fciptunnel create** allow you to control specific features. Use the [O|I] options only with **fciptunnel modify**.
- ficon-tera-read [*disable* | *enable*]**
Enables or disables FICON Read Emulation for a Teradata server on the specified extension tunnel.
- ficon-tera-write [*disable* | *enable*]**
Enables or disables FICON Write Emulation for a Teradata server on the specified extension tunnel.
- ficon-xrc [*disable* | *enable*]**
Enables or disables FICON XRC emulation. FICON XRC Emulation allows XRC (IBM eXtendedRemote Copy, also known as IBM z/OS Global Mirroring) to operate effectively at extended distances.
- ficon-tape-write [*disable* | *enable*]**
Enables or disables FICON Tape Write Pipelining. This feature improves the performance of certain applications when writing to tape over extended distances.

- ficon-tape-read [disable | enable]**
Enables or disables FICON Tape Read Pipelining. This feature improves performance for certain applications when reading from FICON tape over extended distances.
- ficon-tin-tir [disable | enable]**
Enables or disables FICON TIN/TIR emulation. This feature enhances recovery when a TIN/TIR exchange occurs as part of a channel recovery operation during tape emulation. This feature is enabled by default (recommended setting).
- ficon-dvckack [disable | enable]**
Enables or disables FICON Device Level Acknowledgement emulation. This feature is applicable to both FICON Disk and Tape configurations. The feature removes one network round trip for exchanges that end with a Device Level Acknowledgement frame from the device. This feature is enabled by default (recommended setting).
- ficon-read-blk [disable | enable]**
Enables or disables FICON read Tape Read Block ID emulation. This feature permits FICON write channel programs containing embedded read block ID commands (CCWs) with a byte count of exactly four bytes to be processed as emulated commands during write emulation processes.
- ficon-print [disable | enable]**
Enables or disables FICON printer emulation on the specified FCIP tunnel. This command is valid only with the **modify** option.
- max-read-pipe *value***
Defines the maximum number of tape read channel commands (CCWs) that can enter the read pipeline for a single device whether all the CCWs are bundled in a single channel program or in multiple channel programs. The setting has significance only for host (channel) initiated operations at this side and will not affect tape write operations initiated by hosts (channels) attached at the opposite side. Too small of a value will result in poor performance. The value should be chosen based upon the typical tape channel program that requires optimum performance. The default value is 32 (recommended setting). The range is 1 to 100.
- max-write-pipe *value***
Defines the maximum number of tape write channel commands (CCWs) that can enter the write pipeline for a single device whether all the CCWs are bundled in a single channel program or in multiple channel programs. The setting has significance only for host (channel) initiated operations at this side and will not affect tape write operations initiated by hosts (channels) attached at the opposite side. Too small of a value will result in poor performance. The value should be chosen based upon the typical tape channel program that requires optimum performance. The default value is 32 (recommended setting). The range is 1 to 100.
- max-read-devs *value***
Defines the maximum number of concurrent emulated tape read operations. As concurrency increases, the value of emulation decreases. Excessive concurrency has the potential to oversubscribe packet data memory. The setting has significance only for host (channel) initiated operations at this side and will not affect tape read operations initiated by hosts (channels) attached at the opposite side. The default value is 16. The range is 1 to 32.

--max-write-devs *value*

Defines the maximum number of concurrent emulated tape write operations. As concurrency increases, the value of emulation decreases. Excessive concurrency has the potential to oversubscribe packet data memory. The setting has significance only for host (channel) initiated operations at this side and will not affect tape write operations initiated by hosts (channels) attached. The default value is 16. The range is 1 to 32.

--write-timer *value*

Defines a time limit for pipelined write chains. This value is specified in milliseconds (ms). If a pipelined write chain takes longer than this value to complete, the ending status for the next write chain will be withheld from the channel. This limits processing to what the network and device can support. Too small a value limits pipelining performance. Too large a value results in too much data being accepted for one device on a path. The default value is 300 milliseconds (ms). The range is 100 to 1500.

--write-chain *value*

Defines the maximum amount of data that can be contained in a single CCW chain. If this value is exceeded, emulation is suspended. The default value is 3.2 MB (3200000 bytes). The range is 1 MB to 5 MB.

--oxid-base *value*

Defines the base value of an entry pool of 256 OXIDs supplied to emulation-generated exchanges. It should fall outside the range used by FICON channels and devices to avoid conflicts. The default value is 0x8000 (recommended setting). The range is 0x0000 to 0xF000.

--ficon-debug *value*

Defines optional debug flags. The default value is 0xF7C80000. This parameter is primarily for use by technical support personnel.

Optional FCIP circuit arguments for **fciptunnel create** and **modify** include the following.

-a | --admin-status [disable | enable]

Enables or disables the circuit. Admin status is enabled by default.

-S | --local-ip *ipaddress* | none

Sets the local IP address to use for the circuit. The **none** option is supported only on the Brocade 7840 switch and the Brocade SX6 extension blade.

-D | --remote-ip *ipaddress* | none

Sets the remote IP address to use for the circuit. The **none** option is supported only on the Brocade 7840 switch and the Brocade SX6 extension blade.

--local-ha-ip *ipaddress* | none

Sets the local HA IP address to use for the circuit. This argument is applicable on the Brocade 7840 switch and the Brocade SX6 extension blade only.

--remote-ha-ip *ipaddress* | none

Sets the remote HA IP address to use for the circuit. This argument is applicable on the Brocade 7840 switch and the Brocade SX6 extension blade only.

- C | --connection-type** *[default | listener | initiator]*
 Specifies whether the circuit is the listener or the initiator. In default mode, the initiator and listener are automatically chosen based on the lower and higher-order IP address. This can cause a problem in Network Address Translation (NAT) environments, if both sides of the circuit have lower-order addresses.
- L | --load-leveling** *[default | failover | spillover]*
 Configures spillover or failover load-balancing method. The default load-balancing method is failover. This argument is applicable on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- s | --sack** *[disable | enable]*
 Disables or enables selective acknowledgement code (SACK) on the extension circuit. SACK is enabled by default. Use **-s | --sack** to disable the feature when you create a circuit. Use **-s | --sack disable | enable** to disable or enable SACK when you modify a circuit. This operand is supported on the Brocade FX8-24 blade only.
- k | --keepalive-timeout** *timeout*
 Specifies the keep alive timeout in milliseconds. The valid range is 500 ms to 7200000 ms. If the tunnel does not already have FICON Emulation enabled, circuits created on the tunnel default to the the keep alive timeout of 10000 ms (10 seconds) for Brocade FX8-24 blade and 6000 ms (6 seconds) for Brocade 7840 switch and the Brocade SX6 extension blade. If FICON emulation is enabled on the extension tunnel when a circuit is created, the keep alive timeout defaults to 1000 ms (1 seconds).
- x | --metric** *metric*
 Specifies the metric for the configured circuit. The valid range is 0 to 1. The default value is 0. A lower metric assigns a higher priority to the circuit. As data is flowing through the extension tunnel, it automatically traverses the lowest metric cost circuits. For example, if a tunnel has four circuits, three of which are set to a metric of 0 and one is set to a metric of 1, all data will flow over the metric 0 circuits. This parameter is meaningful only, if you configure more than one circuit.
- g | --failover-group** *failover_group_ID*
 Specifies the failover group ID for the configured circuit. The valid range is 0 to 9, where 0 is the default failover group. The circuit failover groups must be defined at both ends of the extension tunnel and each failover group should include at least one metric 0 and one metric 1 circuit. If all metric 0 circuits in the failover group go down, the transmission fails over to the metric 1 circuits in the group. If all metric 0 circuits in a tunnel go down, by default the traffic will run over the metric 1 circuits.
- m | --min-retrans-time** *time*
 Specifies the minimum time interval in milliseconds between retransmits. The valid range is 20 ms to 5000 ms. The default value is 100 ms. This operand is supported on the Brocade FX8-24 blade only.
- r | --max-retransmits** *retransmissions*
 Specifies the maximum number of retransmissions. The valid range is 1 to 16. The default value is 8. This operand is supported on the Brocade FX8-24 blade only.

-v | --vlan-tagging *vlan_id*

Creates an extension tunnel with VLAN Tagging and Class of Service (CoS). Specify a *vlan_id* in the range between 1 and 4094. If any other VLAN option is specified, the VLAN ID must also be specified. Refer to the IEEE 802.1p specification for more information.

You can configure VLAN tags when you create a circuit or after the fact by modifying a circuit. Note that adding or modifying the VLAN configuration after a circuit has been created is a disruptive operation. This operand is supported on the Brocade FX8-24 blade only.

You may also specify VLAN tagging per IP interface with the **portCfg vlantag** command. The VLAN configuration at the IP interface level is for non-data path traffic only. If the data path traffic is to be tagged, it must be done through the VLAN tagging option with the **fcipcircuit create** or **modify** command. Note that the circuit VLAN configuration takes priority over the IP interface VLAN configuration.

--l2cos-f-class *l2cos*

Specifies the Layer 2 Class of Service (L2CoS) value for F-Class Traffic. This priority setting controls connections between switches. The range is 0 to 7. The default is 0.

--l2cos-high *l2cos*

Specifies the L2CoS value for FC High Priority Traffic. The range is 0 to 7. The default is 0.

--l2cos-medium *l2cos*

Specifies the L2CoS value for FC Medium Priority Traffic. The range is 0 to 7. The default is 0.

--l2cos-low *l2cos*

Specifies the L2CoS value for FC Low Priority Traffic. The range is 0 to 7. The default is 0.

--l2cos-ip-high *l2cos*

Specifies the L2CoS value for IP High Priority Traffic. The range is 0 to 7. The default is 0. This parameter is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.

--l2cos-ip-medium *l2cos*

Specifies the L2CoS value for IP Medium Priority Traffic. The range is 0 to 7. The default is 0. This parameter is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.

--l2cos-ip-low *l2cos*

Specifies the L2CoS value for IP Low Priority Traffic. The range is 0 to 7. The default is 0. This parameter is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.

--dscp-f-class *dscp*

Specifies the DSCP value for F-Class Traffic. The range is 0 to 63. The default value is 0.

- dscp-high *dscp***
Specifies the DSCP value for FC High Priority. The range is 0 to 63. The default value is 0.
- dscp-medium *dscp***
Specifies the DSCP value for FC Medium Priority. The range is 0 to 63. The default value is 0.
- dscp-low *dscp***
Specifies the DSCP value for FC Low Priority. The range is 0 to 63. The default value is 0.
- dscp-ip-high *dscp***
Specifies the DSCP value for IP High Priority. The range is 0 to 63. The default value is 0. This parameter is supported only on Brocade 7840 switch and the Brocade SX6 extension blade.
- dscp-ip-medium *dscp***
Specifies the DSCP value for IP Medium Priority. The range is 0 to 63. The default value is 0. This parameter is supported only on Brocade 7840 switch and the Brocade SX6 extension blade.
- dscp-ip-low *dscp***
Specifies the DSCP value for IP Low Priority. The range is 0 to 63. The default value is 0. This parameter is supported only on Brocade 7840 switch and the Brocade SX6 extension blade.

delete *tunnel_ID*

Deletes the specified extension tunnel. This command deletes all associated circuits created with the **fciptunnel** or **fcipcircuit** commands. Use the **portShow** command to display all tunnels and their associated circuits.

fcipcircuit

Creates an extension circuit on an existing tunnel. Use this command to configure additional circuits. The circuit-specific parameters are optional. The syntax for **portcfg fcipcircuit** is as follows:

```
portcfg fcipcircuit [slot]/ve_port option circuit_ID options [arguments] [optional_arguments]
```

The following options and arguments are supported with **fcipcircuit**:

```
create circuit_ID -D | --remote-ip remote_ip_addr -S | --local-ip local_ip_addr -b | --min-comm-rate value in kbps -B | --max-comm-rate value in kbps [circuit_arguments]
```

Creates an extension circuit. You must specify the following parameters when creating an additional circuit:

circuit_ID

Specifies a numeric identifier for the circuit. The circuit ID is an integer value between 0-9 for both the Brocade FX8-24, and 0-7 for the Brocade 7840 switch.

-D | **--remote-ip** *remote_ip_addr*

Specifies the IP address for the remote end of the circuit.

--local-ha-ip *ipaddress* | none

Sets the local HA IP address to use for the circuit. This argument is applicable on the Brocade 7840 switch and the Brocade SX6 extension blade only.

--remote-ha-ip *ipaddress* | none

Sets the remote HA IP address to use for the circuit. This argument is applicable on the Brocade 7840 switch and the Brocade SX6 extension blade only.

-S | --local-ip *local_ip_addr*

Specifies the IP address for the local end of the circuit.

-b | --min-comm-rate *value in kbps* -B | --max-comm-rate *value in kbps*

You can set a minimum and a maximum for the committed rate to configure the tunnel for Adaptive Rate Limiting (ARL), which allows for a more effective sharing of bandwidth between applications. For Brocade FX8-24, the valid range is 10,000 Kbps to 1,000,000 Kbps for 1G ports and 10,000 Kbps to 10,000,000 Kbps for 10G ports. For Brocade 7840 and the Brocade SX6 extension blade, the valid range is 20,000 Kbps to 10,000,000 Kbps. The maximum committed rate can be no larger than five times the minimum committed rate, and both sides of the circuit must have matching configurations.

modify [*circuit_ID*] [*circuit_arguments*]

Modifies the properties of an existing extension circuit. To modify a circuit, you must specify at least one of the optional circuit parameters for the command to be effective. Any circuit attribute you change with the **fcipcircuit modify** command affects only the specified circuit. All other circuits remain unchanged.

--sla *sla_name* | none

Assigns an SLA to a circuit. Make sure to configure the other end of the circuit with a matching SLA. Specify **none** to remove the SLA for a circuit.

Refer to **fciptunnel create** and **modify** for a listing of other optional circuit arguments and their descriptions.

delete *circuit_ID*

Deletes the specified extension circuit.

Examples

To create a FICON-enabled tunnel, first create an empty extension tunnel without any circuits:

```
switch:admin> portcfg fciptunnel 1/12 create
Operation succeeded
```

To enable FICON and compression on the tunnel (this sets the keepalive timeout value to 1000 ms, which is the default for FICON-enabled tunnels):

```
switch:admin> portcfg fciptunnel 1/12 modify -c hardware -F enable
Operation succeeded
```

To create two circuits (circuit 0 and circuit 1) on the tunnel (these circuits will be created with the FICON-compatible keepalive timeout value):

```
switch:admin> portcfg fcipcircuit 1/12 create \
0 192.168.51.61 192.168.50.68 -b 300000 -B 500000
Operation succeeded
```

```
switch:admin> portcfg fcipcircuit 1/12 create
1 192.168.50.61 192.168.51.68 -b 300000 -B 500000
Operation succeeded
```

To display circuit 1:

```
switch:admin> portshow fcipcircuit 1/12 1
-----
Circuit ID: 1/12.1
  Circuit Num: 1
  Admin Status: Enabled
  Oper Status: In Progress
  Remote IP: 192.168.51.61
  Local IP: 192.168.51.68
  Metric: 0
  Failover Group ID: (Not Config/Active)
  Min Comm Rt: 300000
  Max Comm Rt: 500000
  SACK: On
  Min Retrans Time: 100
  Max Retransmits: 8
  Keepalive Timeout: 1000
  Path MTU Disc: 0
  VLAN ID: (Not Configured)
  L2CoS: (VLAN Not Configured)
  DSCP: F: 0 H: 0 M: 0 L: 0
  Flags: 0x00000000
```

To set the compression rate to 'moderate' on the tunnel:

```
switch:admin> portcfg fciptunnel 1/12 modify -c moderate
Operation succeeded
```

To enable FICON XRC and Teradata emulation on the tunnel:

```
switch:admin> portcfg fciptunnel 1/12 modify \
--ficon-xrc enable --ficon-tera-read enable --ficon-tera-write enable
Operation succeeded
```

To enable printer emulation on the tunnel:

```
switch:admin> portcfg fciptunnel 1/12 modify \
--ficon-print enable
Operation succeeded
```

To configure FCIP FastWrite and Tape Pipelining on the tunnel:

```
switch:admin> portcfg fciptunnel 1/12 modify -f enable -t enable
Operation succeeded
```

To set the bandwidth distribution mode to "protocol" and adjust the FC and IP bandwidth ratios for an IP-Extension enabled tunnel on a Brocade 7840 switch:

```
switch:admin> portcfg fciptunnel 24 create --distribution protocol,60,40
Operation succeeded
```

To create a circuit on the tunnel with the failover group ID and verify the configuration (in this example, the OpStatus "FGrpWrn" indicates that the failover group is defined but there is not at least one metric 0 and one metric 1 circuit as part of the failover group):

```
switch:admin> portcfg fcipcircuit 1/22 create 2 \
--remote-ip 1.42.128.93 --local-ip 1.42.128.23 --min-comm-rate 500000 --max-comm-rate 500000
-x 1-g 1
Operation succeeded
switch:admin> portshow fciptunnel all -c
```

| Tunnel | Circuit | OpStatus | Flags | Uptime | TxMBps | RxMBps | ConnCnt | CommRt | Met/G |
|--------|----------|----------|---------|--------|--------|--------|---------|-----------|-------|
| 1/22 | - | FGrpWrn | cft---- | 20m26s | 0.00 | 0.00 | 1 | - | -/- |
| 1/22 | 0 1/xge0 | Up | ---4--s | 20m26s | 0.00 | 0.00 | 1 | 3000/5000 | 0/- |
| 1/22 | 1 1/xge0 | Up | ---4--s | 3s | 0.00 | 0.00 | 2 | 3000/5000 | 0/- |
| 1/22 | 2 1/xge0 | Up | ---4--s | 2m7s | 0.00 | 0.00 | 1 | 5000/5000 | 1/1 |

To set the failover group for circuit 1 and verify the configuration:

```
switch:admin> portcfg fcipcircuit 1/12 modify 1-g 1
Operation succeeded
switch:admin> portshow fciptunnel all -c
```

| Tunnel | Circuit | OpStatus | Flags | Uptime | TxMBps | RxMBps | ConnCnt | CommRt | Met/G |
|--------|----------|----------|---------|--------|--------|--------|---------|-----------|-------|
| 1/22 | - | Up | cft---- | 26m51s | 0.00 | 0.00 | 1 | - | -/- |
| 1/22 | 0 1/xge0 | Up | ---4--s | 20m26s | 0.00 | 0.00 | 1 | 3000/5000 | 0/- |
| 1/22 | 1 1/xge0 | Up | ---4--s | 3s | 0.00 | 0.00 | 2 | 3000/5000 | 0/1 |
| 1/22 | 2 1/xge0 | Up | ---4--s | 2m7s | 0.00 | 0.00 | 1 | 5000/5000 | 1/1 |

To delete circuit 1:

```
switch:admin> portcfg fcipcircuit 1/12 delete 1
Operation succeeded
```

To create a tunnel with VLAN tagging:

1. Create a tunnel with VLAN tagging set.

```
switch:admin> portcfg fciptunnel 16 create \
-D 192.168.2.10 -S 192.168.2.10 -b 2500000 -B 2500000 -v 100
Operation Succeeded
```

2. Create a circuit with VLAN tagging set.

```
switch:admin> portcfg fcipcircuit 16 create 1 \
-D 192.168.2.11 -S 192.168.2.11 -b 2500000 -B 2500000 -v 200
Operation Succeeded
```

3. Modify an existing circuit to change the VLAN tag and L2 CoS levels.

```
switch:admin> portcfg fcipcircuit 16 modify \
O -v 300 --l2cos-f 7 --l2cos-h 5 \
--l2cos-m 3 --l2cos-l 1
!!!! WARNING !!!!
Modify operation can disrupt the traffic on the
fcipcircuit specified for a brief period of time. This
```

```
operation will bring the existing circuit down
(if circuit is up) before applying new configuration.
```

```
Continue with Modification (Y,y,N,n): [n] y
Operation Succeeded
```

4. Modify existing circuit to change DSCP marking values

```
switch:admin> portcfg fcipcircuit 16 modify 0 \
--dscp-f 32 --dscp-h 16 --dscp-m 8 --dscp-l 4
Operation Succeeded
```

5. Display the tunnel configuration and circuit values.

```
switch:admin> portshow fciptunnel 24 -c
-----
Tunnel: VE-Port:24 (idx:0, DP0)
=====
Oper State           : Online
TID                  : 24
Flags                : 0x00000000
IP-Extension        : Disabled
Compression         : None
QoS BW Ratio        : 50% / 30% / 20%
Fastwrite           : Disabled
Tape Pipelining     : Disabled
IPSec               : Enabled
IPSec-Policy        : myPolicy1
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN           : 10:00:00:05:33:65:82:c8
Peer WWN            : 10:00:00:05:33:65:84:08
RemWWN (config)    : 00:00:00:00:00:00:00:00
cfgmask             : 0x0000001f 0x4000020c
Flow Status         : 0
ConCount/Duration   : 2 / 2h41s
Uptime              : 1h31m59s
Stats Duration      : 1h31m59s
Receiver Stats      : 1020072 bytes / 6534 pkts / 210.00 Bps Avg
Sender Stats        : 835204 bytes / 6535 pkts / 173.00 Bps Avg
TCP Bytes In/Out    : 28609996 / 35324148
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)   : 1 / 1 / 1 ms
Wan Util            : 0.0%
TxQ Util            : 0.0%

Circuit 24.0 (DP0)
=====
Admin/Oper State    : Enabled / Online
Flags               : 0x00000000
IP Addr (L/R)       : 10.1.8.77 ge8 <-> 10.1.8.76
HA IP Addr (L/R)    : 0.0.0.0 ge0 <-> 0.0.0.0
Configured Comm Rates: 1000000 / 1000000 kbps
Peer Comm Rates     : 1000000 / 1000000 kbps
Actual Comm Rates   : 1000000 / 1000000 kbps
Keepalive (Cfg/Peer): 6000 (6000 / 6000) ms
Metric              : 0
Connection Type     : Default
ARL-Type            : Auto
PMTU                : Disabled
SLA                 : (none)
```



```

Failover Group      : 0
VLAN-ID            : NONE
L2Cos (FC:h/m/l)   : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)   : 0 / 0 / 0
DSCP (FC:h/m/l)    : 16 / 8 / 4 (Ctrl:32)
DSCP (IP:h/m/l)    : 0 / 0 / 0
cfgmask            : 0x40000000 0x01e10c2f
Flow Status        : 0
ConCount/Duration  : 2 / 2h42s
Uptime             : 1h31m59s
Stats Duration     : 1h31m59s
Receiver Stats     : 515664 bytes / 3307 pkts / 97.00 Bps Avg
Sender Stats       : 415636 bytes / 3269 pkts / 77.00 Bps Avg
TCP Bytes In/Out   : 19805096 / 23191360
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)  : 1 / 1 / 1 ms
Wan Util           : 0.0%

```

Circuit 24.1 (DP0)

```

=====

```

```

Admin/Oper State   : Enabled / Online
Flags              : 0x00000000
IP Addr (L/R)      : 192.168.0.20 ge6 <-> 192.168.0.10
HA IP Addr (L/R)   : 0.0.0.0 ge0 <-> 0.0.0.0
Configured Comm Rates: 1000000 / 1000000 kbps
Peer Comm Rates    : 1000000 / 1000000 kbps
Actual Comm Rates  : 1000000 / 1000000 kbps
Keepalive (Cfg/Peer) : 6000 (6000 / 6000) ms
Metric             : 0
Connection Type    : Default
ARL-Type           : Auto
PMTU               : Disabled
SLA                : (none)
Failover Group     : 0
VLAN-ID            : NONE
L2Cos (FC:h/m/l)   : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)   : 0 / 0 / 0
DSCP (FC:h/m/l)    : 16 / 8 / 4 (Ctrl:32)
DSCP (IP:h/m/l)    : 0 / 0 / 0
cfgmask            : 0x40000000 0x01e00c2f
Flow Status        : 0
ConCount/Duration  : 2 / 1h48m37s
Uptime             : 1h31m59s
Stats Duration     : 1h31m59s
Receiver Stats     : 504408 bytes / 3227 pkts / 103.00 Bps Avg
Sender Stats       : 419568 bytes / 3266 pkts / 88.00 Bps Avg
TCP Bytes In/Out   : 16794424 / 20191296
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)  : 1 / 1 / 1 ms
Wan Util           : 0.0%

```

To modify an existing circuit so it becomes the listener:

```
switch:admin> portcfg fcipcircuit 16 modify 2 -C 1

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fcipcircuit specified \
for a brief period of time. This operation will bring the existing \
circuit down (if circuit is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [n]y
Operation Succeeded
```

To modify an existing circuit so it becomes the initiator (the following example uses the string option):

```
switch:admin> portcfg fcipcircuit 16 \
              modify 2 --connection-type initiator

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fcipcircuit specified \

for a brief period of time. This operation will bring the existing \
circuit down (if circuit is up) before applying new configuration.

Continue with Modification (Y,y,N,n): [n]y
Operation Succeeded
```

To configure an IPSec-enabled tunnel to connect with an IPSec-configured tunnel in legacy mode:

```
switch:admin> portcfg fciptunnel 8/12 modify --legacy enable

!!!! WARNING !!!!
Modify operation can disrupt the traffic on the fciptunnel specified \
for a brief period of time. This operation will bring the existing \
tunnel down (if tunnel is up) before applying new configuration.\
Continue with Modification (Y,y,N,n): [n]y
Tunnel 8/12 modify: Operation Succeeded
```

SEE ALSO [configure](#), [portCfgGE](#), [portCfgShow](#), [portCmd](#), [portShow](#), [switchShow](#)

portCfgAlpa

Configures the AL_PA offset on a specified port or range of ports.

Synopsis `portcfgalpa [slot/]port, mode`

Description Use this command to set the Arbitrated Loop Physical Address (AL_PA) offset on a port or a range of ports to either 0x0 (default) or 0x13.

Changes made by this command are persistent across switch reboots and power cycles.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands When invoked without operands, this command displays the usage. The following operands are supported:

slot

For bladed systems only, specify the slot number of the port to be configured, followed by a slash (/).

port

Specify the number of the port to be configured, relative to its slot for bladed systems. Use the **switchShow** command for a list of valid ports.

mode

Specify a value of 1 to set the AL_PA to 0x13. A value of 0 sets the default AL_PA to 0x0. This operand is required.

Examples To configure a port with AL_PA 0x0 (default):

```
switch:admin> portcfgalpa 1/3 0
```

To configure a port with AL_PA 0x13:

```
switch:admin> portcfgalpa 1/3 1
```

SEE ALSO [portCfgShow](#)

portCfgAutoDisable

Manages the port autodisable configuration.

Synopsis

```
portcfgautodisable --enable [slot/]port[-port]
portcfgautodisable --disable [slot/]port[-port]
portcfgautodisable --add option [slot/]port[-port]
portcfgautodisable --remove option [slot/]port[-port]
portcfgautodisable --addall [slot/]port[-port]
portcfgautodisable --removeall [slot/]port[-port]
portcfgautodisable --addexcept option [slot/]port[-port]
portcfgautodisable --show [slot/]port[-port]
portcfgautodisable --suspend
portcfgautodisable --resume
portcfgautodisable --help
```

Description Use this command to enable or disable the autodisable feature for a specified port or a range of ports and to manage the configuration. If the ports are already in the requested configuration, no action is taken. If a range of ports is specified, some of which are already in the requested configuration, no action is taken for those ports. All other ports in the specified range are updated. Execution of this command is nondisruptive.

The autodisable feature is by default disabled for all ports.

Use the **--suspend** and **--resume** options respectively to temporarily suspend and resume the configuration in the current logical switch. These states are persistent across reboots and HA failover.

The port autodisable feature minimizes traffic disruption introduced in some instances of automatic port recovery. When the autodisable flag is set, you can specify the conditions that will prevent the port to reinitialize. Such conditions include loss of sync, loss of signal, OLS, NOS, and LIP. Refer to the operand section for an explanation of these conditions. Note that a link reset does not cause a port autodisable. When a port is in FICON Management Server (FMS) mode, an autodisable port remains persistently disabled across High Availability (HA) failover. In all cases, you can bring the automatically disabled port back into service using the **portEnable** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on GbE ports, VE_Ports, logical interswitch links (LISLs), or FCoE ports.

Operands This command has the following operands:

slot Specifies the slot number on bladed systems, followed by a slash (/).

| | |
|--------------------|--|
| <i>port[-port]</i> | Specifies a port or a port range, relative to the slot number on bladed systems, for example, 5/17-29. |
| --enable | Enables the autodisable feature on the specified ports. |
| --disable | Disables the autodisable feature on the specified ports. |
| --add | Specifies one or more trigger conditions that will disable the specified ports. Trigger conditions must be separated by a space, for example, LOSN OLS LIP. Trigger conditions are case insensitive. |
| --remove | Removes one or more trigger conditions from the port autodisable configuration. |
| --addall | Adds all supported trigger conditions to the port autodisable configuration. |
| --removeall | Removes all supported trigger conditions from the port autodisable configuration. This command removes the triggers only; it does not disable the port autodisable flag. |
| --addexcept | Adds all supported trigger conditions to the port autodisable configuration except the one specified . |
| --show | Displays the port autodisable configuration. |
| --suspend | Suspends the current port autodisable configuration temporarily. |
| --resume | Resumes the suspended port autodisable configuration. |
| --help | Displays command usage. |
| <i>option</i> | Specify one or more the following options (supported with the --add , --remove , and --addexcept operands): |
| LOSN | Loss of Sync |
| LOSNG | Loss of Signal |
| OLS | Offline Primitive Sequence |
| NOS | Not Operational Primitive Sequence |

LIP

Loop Initialization Primitive Sequence

Examples To disable the port autodisable feature on single port:

```
switch:admin> portcfgautodisable --disable 18
switch:admin> portcfgautodisable --show 18
Port Auto Disable: OFF
Configured Option(s):No events configured. Port will
not be automatically disabled.
```

To enable the port autodisable feature on a single port:

```
switch:admin> portcfgautodisable --enable 18
switch:admin> portcfgautodisable --show 18
Port Auto Disable: ON
MODE                : RESUME
Configured Option(s):No events configured. Port will
not be automatically disabled.
```

To add multiple trigger conditions to the port autodisable configuration on a single port:

```
switch:admin> portcfgautodisable --add losn lip ols 18
switch:admin> portcfgautodisable --show 18
Port Auto Disable: ON
MODE                : RESUME
Configured Option(s): LOSN OLS LIP
```

To add all trigger conditions except the LOSN condition:

```
switch:admin> portcfgautodisable --addexcept LOSN 18
switch:admin> portcfgautodisable --show 18
Port Auto Disable: ON
MODE                : RESUME
Configured Option(s): LOSG OLS NOS LIP
```

To attempt to enable and configure the port autodisable feature on a port for which the feature is already activated (the configuration is not updated):

```
switch:admin> portcfgautodisable --show 18
Port Auto Disable: ON
MODE                : RESUME
Configured Option(s): LOSG OLS NOS LIP
switch:admin> portcfgautodisable --enable 18
Same configuration for port 18
switch:admin> portcfgautodisable --addexcept LOSN 18
Same configuration for port 18
```

To enable the port autodisable feature on a range of ports:

```
switch:admin> portcfgautodisable --enable 0-2
switch:admin> portcfgautodisable --show 0-2
Port Auto Disable: ON
MODE                : RESUME
Configured Option(s):No events configured.
Port will not be automatically disabled.
```

```
Port Auto Disable: ON
```

```

MODE                : RESUME
Configured Option(s):No events configured.
Port will not be automatically disabled.

```

```

Port Auto Disable: ON
MODE                : RESUME
Configured Option(s):No events configured.
Port will not be automatically disabled.

```

To enable the port autodisable feature for a range of ports, some of which were previously enabled (the configuration is applied only to those ports, for which an update is necessary, that is, port 3 in the following example):

```

switch:admin> portcfgautodisable --enable 0-3
Same configuration for port 0
Same configuration for port 1
Same configuration for port 2

```

To suspend the port autodisable configurations:

```

switch:admin> portcfgautodisable --suspend
Suspending Port Auto Disable will de-activate the triggers \
  for all ports configured with Port autodisable conditions on this Logical Switch.
Would you like to continue [y/n]?: y
PAD option is set to 'SUSPEND'
switch:admin> portcfgautodisable --show 18
Port Auto Disable: ON
MODE                : SUSPEND
Configured Option(s): LOSG  OLS  NOS  LIP

```

To resume the suspended port autodisable configurations:

```

switch:admin> portcfgautodisable --resume
PAD option is set to 'RESUME'
switch:admin> portcfgautodisable --show 18
Port Auto Disable: ON
MODE                : RESUME
Configured Option(s): LOSG  OLS  NOS  LIP

```

SEE ALSO [portCfgShow](#), [portEnable](#)

portCfgCompress

Configures a port for compression.

Synopsis

```
portcfgcompress --enable [slot]/port
portcfgcompress --disable [slot]/port
portcfgcompress --help
```

Description Use this command to configure a port for compression. This command enables or disables the compression configuration on the specified port, and saves the configuration persistently.

Configuring a port for compression is disruptive. You must disable the port before you can enable compression on the port. This command fails on an enabled port.

The number of configurable ports is limited per ASIC. A message will be displayed once the maximum number is exceeded. Use the **portEncCompShow** for a listing of configurable ports per ASIC.

Notes This command is supported only on E_Ports. An E_Port can be enabled for compression and for encryption at the same time.

When you move a configured port to another logical switch, you are informed that the operation requires the port configuration to be disabled. You are given the choice to cancel the move or to continue. If you want to go ahead with the move and the port is configured for encryption or compression, you must disable the configuration prior to moving the port. You must reconfigure the port on the target switch if you want to use compression or encryption on that port.

This command is supported only on 16 Gbps-capable platforms running Fabric OS v7.0.0 or later with the exception of the Brocade 6505 switch and the Brocade FC8-32E and Brocade FC8-48E port blades.

Compression feature is not supported on 16 Gbps-capable Inter Chassis Link (ICL) ports.

On 16 Gbps-capable platforms, both encryption and compression features are supported. On 32 Gbps-capable platforms, only compression is supported and encryption is not supported.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot On bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port Specifies the port to be configured.

--enable Enables the compression configuration on the specified port.

--disable Disables the compression configuration on the specified port.

--help
Displays the command usage.

Examples To enable the compression configuration:

```
switch:admin> portcfgcompress --enable 2

switch:admin> lscfg --config 15 -port 2
This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Checking for cap port 2
After Checking for cap port 2
val 1 port 2
lscfg: The port cannot be moved to the requested switch because
port or ports have Encryption/Compression enabled. Please disable
the Encryption/Compression configuration.
  FID:      15
  Slot:     0
  Port:     2
```

To disable the compression configuration:

```
switch:admin> portcfgcompress --disable 2
```

To enable compression configuration when maximum number is reached:

```
switch:admin> portcfgcompress --enable 2
Configuration is not allowed. Maximum number of ports is
already configured for Encryption/Compression.
```

SEE ALSO [portCfgEncrypt](#), [portEncCompShow](#)

portCfgCreditRecovery

Enables or disables credit recovery on a port.

Synopsis `portcfgcreditrecovery --enable [slot]port`

`portcfgcreditrecovery --disable [slot]port`

`portcfgcreditrecovery --help [slot]port`

Description Use this command to enable or disable credit recovery on a port.

The credit recovery feature enables credits or frames to be recovered. The default credit recovery configuration is enabled.

Notes This command is supported in Access Gateway mode.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands This command has the following operands:

--disable Disables credit recovery configuration on the specified port.

--enable Enables credit recovery configuration on the specified port.

--help Displays the command usage.

Examples To enable credit recovery on a port:

```
switch:admin> portcfgcreditrecovery --enable 3/15
```

To disable credit recovery on a port:

```
portcfgcreditrecovery --disable 3/15
```

SEE ALSO [portCfgShow](#)

portCfgDefault

Resets the port configuration to factory default value.

Synopsis `portcfgdefault [slot][ge]port`

`portcfgdefault --help`

Description Use this command to reset all configuration values on a specified port to their factory defaults. This command persistently disables ports capable of routing, which is the factory default value. Use the **portCfgShow** command to display the port configuration.

This command does not change the state of a port. To change the state of an E_Port, use either **switchDisable/switchEnable** or **portDisable/portEnable**.

When this command is used to reset an F_Port in an Access Gateway, that F_Port is unmapped from its mapping to an N_Port.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be reset, followed by a slash (/).

port

Specifies the number of the port to be reset, relative to its slot for bladed systems. Use **switchShow** to list of valid ports.

`--help`

Displays the command usage.

Examples To reset a port to factory defaults:

```
switch:admin> portcfgdefault 1/3
```

To reset a GbE port to factory defaults:

```
switch:admin> portcfgdefault 8/ge1
```

SEE ALSO [portCfgEport](#), [portCfgGport](#), [portCfgLongDistance](#), [portCfgLport](#), [portCfgPersistentDisable](#), [portCfgPersistentEnable](#), [portCfgShow](#), [portCfgSpeed](#), [portCfgTrunkPort](#)

portCfgDPort

Configures a port as a D_Port.

Synopsis

```
portcfgdport --enable [-dwdm] [slot/]port_list
portcfgdport --disable [-dwdm] [slot/]port_list
portcfgdport --provision [-add] [-dwdm] [slot/]port_list
portcfgdport --provision [-delete] [-dwdm] [slot/]port_list
portcfgdport --provision [-show] [slot/]port_list
portcfgdport --help
```

Description Use this command to configure a diagnostic port (D_Port). The D_Port is not part of the fabric. It does not carry any interswitch traffic or data traffic. The D_Port is used solely for the purpose of running link-level diagnostics between two switches, switch-HBA, AG-AG, or AG-HBA and to isolate link level fault on the port, in the small form factor pluggable (SFP), or in the cable.

You must configure both ends of the link between a given pair of switches, and you must disable the port before you can configure a D_Port. Re-enabling the D_Ports will automatically start the diagnostics when the ports come online.

The D_Port test performs the following diagnostics:

- An electrical loopback test (supported only on 16 Gb and 32 Gb SFPs capable of electrical loopback)
- An optical loopback test (supported only on 16 Gb and 32 Gb SFPs capable of optical loopback)
- A link traffic test
- A link distance measurement

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported only on Fibre Channel ports. SFPs must be 10 Gb or 16 Gb Brocade-branded SFPs and run Fabric OS v7.0.0 or later on 16 Gb-capable platforms.

Beginning with Fabric OS v7.2.1 or later, the command is also supported on 8 Gb LWL/ELWL Brocade-branded SFPs, QSFPs or QSFP+ on 16 Gb-capable platforms.

D_Port configuration is not supported on EX_Ports, SIM-Ports, and ports configured in encryption mode or compression mode.

Links with mismatched D_Port configuration will be segmented or disabled.

Forward Error Correction (FEC) is not supported on D_Ports configured with Dense Wavelength Division Multiplexing (DWDM).

On-Demand D_Port test must not be run on the links having DWDM as it can cause D_Port test to stuck in progress. If DWDM is present in between, configure the static D-Port test with **-dwdm** option.

Do not run D_Port tests on more than eight links at a time for normal (not long distance) links on FC16-64 blade. On a FC 16-64 blade, do not enable the On-Demand D_Port with **bladeInsert** or **slotPowerOn**.

Run D_Port test on one link at a time for 100 Km long distance ports on FC 16-64 blade. D_Port tests on other long distance links must be started only after the D_Port test completes on the previous long distance link.

Operands

This command has the following operands:

slot

On bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port_list

Specifies one or more ports, relative to the slot on bladed systems. Use **switchShow** for a listing of valid ports. A port list should be enclosed in double quotation marks and can consist of the following:

- A single port, for example, "8" or "5/8" on blades systems.
- A port range where beginning and end port are separated by a dash, for example, "8-13" or "5/8-13" on blades systems. A port range cannot span multiple slots.
- A set of ports separated by commas, for example "3,5,7,8" or "5/3,5,7,8" on bladed systems.
- A wildcard * indicates all ports. The wildcard can be represented as "*" or "***".

--enable

Configures the specified port(s) as D_Port. You cannot configure the port as a D_Port while Dynamic D_Port tests are in progress on the port. The operation fails with an error message.

-dwdm

Skips the optical loopback test when you run the D_Port diagnostics over a Dense Wavelength Division Multiplexing (DWDM) link or any third party equipment that connects two switches. This operand is optional.

--disable

Clears the D_Port configuration on the specified ports.

-dwdm

Disables DWDM mode on the port. This operand is optional.

--provision

List of ports that are marked to be set as D_Ports. Use the **--enable** command to configure one or more ports from the list as D_Ports.

-add

Adds the ports on the provision list to be configured as D_Ports.

-dwdm

Provisions the specified ports for DWDM. This operand is optional.

-delete

Removes the ports from the provision list.

-dwdm

Disables DWDM provisioning on the specified ports. This operand is optional.

-show

Lists all the ports in the provision list.

If the port range is specified, the status of each port on the provision list displays. Displays ON if the port is in the provision list or OFF if the port is not in the provision list.

--help

Displays the command usage.

Examples

To configure a single port as a D_Port:

```
switch:admin> portdisable 42
```

```
switch:admin> portcfgdport --enable 42
```

Caution: D_Port functionality is only available on 16Gb-capable platforms \ with 16Gb FC SFPs, 10Gb FC SFPs, or 8Gb LWL/ELWL FC SFPs, QSFPs or QSFP+.

```
switch:admin> portenable 42
```

To clear the D_Port configuration:

```
switch:admin> portdisable 42
```

```
switch:admin> portcfgdport --disable 42
```

```
switch:admin> portenable 42
```

To add a port to the D_Port provision list:

```
switch:admin> portcfgdport --provision -add 7/16
```

Caution: D_Port functionality is only available on 16Gb-capable platforms \ with 16Gb FC SFPs, 10Gb FC SFPs, 8Gb LWL/ELWL FC SFPs, QSFPs or QSFP+.

To add a port to the D_Port provision list and enable DWDM mode on the port:

```
switch:admin> portcfgdport --provision -add -dwdm 7/16
```

Caution: D_Port functionality is only available on 16Gb-capable platforms \ with 16Gb FC SFPs, 10Gb FC SFPs or 8Gb LWL/ELWL FC SFPs.

To remove a port from the provision list:

```
switch:admin> portcfgdport --provision -delete 4/12
```

To list all the ports in the provision list:

```
switch:admin> portcfgdport --provision -show
```

```
Slot Port D-Port provision DWDM
```

```
=====
```

```
7 16 ON ON
7 17 ON OFF
7 18 ON ON
```

SEE ALSO

[fabricLog](#), [portCfg](#), [portDPortTest](#), [portShow](#), [switchShow](#)

portCfgEncrypt

Configures a port for encryption.

Synopsis `portcfgencrypt --enable [slot]/port`

`portcfgencrypt --disable [slot]/port`

`portcfgencrypt --help`

Description Use this command to configure a port for encryption. This command enables or disables the encryption configuration on the specified port and saves the configuration persistently.

Before you can configure a port for encryption, you must configure the port for authentication. When disabling encryption, you must disable the encryption configuration before you can disable authentication. Refer to the *Fabric OS Encryption Administrator's Guide* for more information.

Configuring a port for encryption is disruptive. You must disable the port before you can enable encryption on the port. This command fails on an enabled port.

The number of configurable ports is limited per ASIC. Use the `portEncCompShow` for a listing of configurable ports per ASIC.

Notes This command is supported on E_Ports and EX_Ports. An E_Port or EX_Port can be enabled for compression and for encryption at the same time.

When you move a configured port to another logical switch, you are informed that the operation requires that the port configuration be disabled. You are given the choice to cancel the move or to continue. If you want to go ahead with the move and the port is configured for encryption or compression, you must disable the configuration prior to moving the port. You must reconfigure the port on the target switch if you want to use compression or encryption on that port.

This command is supported only on 16 Gbps-capable platforms running Fabric OS v7.0.0 or later with the exception of the Brocade 6505 switch and the Brocade FC8-32E and Brocade FC8-48E port blades.

Encryption feature is not supported on 16 Gbps-capable Inter Chassis Link (ICL) ports.

On 16 Gbps-capable platforms, both encryption and compression features are supported. On 32 Gbps-capable platforms, only compression is supported and encryption is not supported.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

On bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port

Specifies the port to be configured.

portCfgEncrypt

- enable** Enables the encryption configuration on the specified port.
- disable** Disables the encryption configuration on the specified port.
- help** Displays the command usage.

Examples

To enable the encryption configuration:

```
switch:admin> portcfgencrypt --enable 2
```

To attempt to move a configured port to another logical switch:

```
switch:admin> lscfg --config 15 -port 2
This operation requires that the affected ports be disabled.
Would you like to continue [y/n]?: y
Checking for cap port 2
After Checking for cap port 2
val 1 port 2
lscfg: The port cannot be moved to the requested switch because
port or ports have Encryption/Compression enabled. Please disable
the Encryption/Compression configuration.
    FID:    15
    Slot:    0
    Port:    2
```

To disable the encryption configuration:

```
switch:admin> portcfgencrypt --disable 2
```

SEE ALSO [authUtil](#), [portCfgCompress](#), [portEncCompShow](#), [secAuthSecret](#)

portCfgEport

Enables or disables E_Port capability on a port or locks down a port as an E_Port.

Synopsis `portcfgeport [slot]port mode`

`portcfgeport -i [index1[-index2] [...] [-f] mode]`

`portcfgeport -slot [slot1[-slot2] [...]`

`portcfgeport -h`

Description Use this command to enable or disable E_Port capability on a port or to lock down a port as an E_Port. E_Port capability is enabled by default. When an interswitch link (ISL) is connected to a port and the port's E_Port capability is disabled, the ISL is segmented, and all traffic between the switches stops. Fabric management data, such as zoning information, can no longer be exchanged through this port.

You can identify a single port to be configured by its port number or by its port index number. Port ranges are supported with index numbers or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers. When used without operands, this command displays all persistently disabled ports on the switch.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **portSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status or alternately use **portSwapShow**.

Changes made by this command are persistent across switch reboots or power cycles.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Regardless of how many E_Ports are connected between two switches, the maximum routing paths are limited to 16 E_Ports.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands This command has the following operands:

| | |
|---------------------------|--|
| <i>slot</i> | For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/). |
| <i>port</i> | Specifies the number of the port to be configured, relative to its slot for bladed systems. Use switchShow to display a listing of valid ports. |
| <i>-i index1[-index2]</i> | Specifies a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, -i 33-38 40-60. |
| <i>-f</i> | Ignores nonexistent ports. This operand is valid only with the -i option. |

portCfgEport

- slot** [*slot1*[-*slot2*]] Specifies all ports on a slot or on a range of slots, for example, **-s 3-5**. You may specify multiple slot ranges separated by a space, for example, **-s 3-5 8-10**.
- mode** Specifies the E_Port configuration as one of the following:
- 0** Disables E_Port capability for the specified ports.
 - 1** Enables the ports as E_Ports. This is the default port state.
 - 2** Locks down the ports as E_Ports. This command effectively disables the port's F_Port capability.
- h** Displays the command usage.

Examples To disable E_Port capability on a port:

```
switch:admin> portcfgeport 1/3 0
```

To enable E_Port capability on a port:

```
switch:admin> portcfgeport 1/3 1
```

To lock down the port as an E_Port.

```
switch:admin> portcfgeport 1/3 2
```

To disable E_Port capability on a range of ports specified by their index number:

```
switch:admin> portcfgeport -i 12-18 0
```

To enable E_Port capability on all ports of slot 3-5:

```
switch:admin> portcfgeport -s 3-5 1
```

SEE ALSO [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portCfgEportCredits

Configures normal distance E_Port buffer allocation.

Synopsis `portcfgeportcredits --enable [slot]port credits`

`portcfgeportcredits --disable [slot]port`

`portcfgeportcredits --show [slot]port`

`portcfgeportcredits --help`

Description

Use this command to change the default credit allocation for a normal distance E_Port or EX_Port by allocating the specified number of credits to the port. When port credit allocation is enabled, the number of credits specified overrides the default E_Port credit allocation; that is, a new credit model is constructed based on the user-configured credits value. When the configuration is disabled, the default credit allocation (default credit model) is restored. Only a normal distance E_Port and EX_Port can utilize the new credit model constructed by this command, and the allocated credits are reserved only for that port.

ICL ports are supported and are used in conjunction with 2 km ICL QSFPs.

The E_Port credit configuration is persistent across system reboots and High Availability (HA) failover.

Use the `portcfgeportcredits --show` command to display the configured credits.

Use the `portBufferShow` command to determine current port buffer allocations.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The E_Port credits feature does not support ports configured as F_Ports, Mirror Ports, L_Ports, and Trunk Areas. If E_Port credits are configured on ports, moving the ports from one logical switch to another logical switch is not permitted.

The Fabric OS port configuration commands are not supported on FCoE ports.

The E_Port credits are mutually exclusive with F_Port buffers, longdistance, ISL R_RDY, and Buffer Optimized mode.

Operands

This command has the following operands:

--enable

Enables the E_Port credit configuration on a specified port. A port and credits allocation must be specified with this option.

--disable

Disables the E_Port credit configuration on a specified port.

--show

Displays the credits configured on a specified port.

slot

Specifies the slot number on bladed systems, followed by a slash (/).

portCfgEportCredits

port

Specifies the port number.

credits

Specifies the number of credits to be allocated to the specified port. The specified credit allocation takes effect when the E_Port comes online. This operand is required with the **--enable** option. The minimum credit allocation is 5 and the maximum is 40 for Gen5 platform (16Gbps-capable blades). For ICL ports of Gen5 platform, the valid range is from 5 through 16. For Gen 6 platform, the credit allocation is allowed in the range between 5 to 160 for both ICL ports and normal ISL and EX Port links. In Gen 6 platform, if the specified credit allocation is less than the default allocation, the default allocation scheme is enforced (in Gen 6 platform, 32 Gbps links are reserved 20 credits per VC). The configured credits will be allocated for each of the medium virtual channels (VCs) for the non-QoS ports. For QoS ports, after sharing, both the medium and high VCs will have the configured credits allocated.

--help

Displays the command usage.

Examples To allocate 10 credits to an E_Port:

```
switch:admin> portcfgeportcredits --enable 12/6 10
Success.
```

To display the configured credits for an E_Port:

```
switch:admin> portcfgeportcredits --show 12/6
E-Port Credit Configured : 10
Success.
```

To disable the credit configuration and return to the default credit allocation:

```
switch:admin> portcfgeportcredits --disable 12/6
Success.
```

SEE ALSO [portBufferShow](#)

portCfgEXPort

Sets a port to be an EX_Port, and sets and displays EX_Port configuration parameters.

Synopsis

```
portcfgexport [slot]port [-port]
portcfgexport [-a admin]
portcfgexport [-f fabricid]
portcfgexport [-r ratov]
portcfgexport [-e edtov]
portcfgexport [-d domainid]
portcfgexport [-p pidformat]
portcfgexport [-t fabric_parameter]
portcfgexport [-m port mode]
portcfgexport [-i mode]
portcfgexport --help
```

Description

Use this command to allow a port to be configured as an EX_Port, to display the port's EX_Port configuration, or to change the configuration. If no optional parameter is specified, the command displays the currently configured values; otherwise, it sets the specified attribute to its new value. The port must be disabled prior to setting EX_Port attributes. The port must be enabled before the port can become active following EX_Port parameter changes. Use **portDisable** and **portEnable** to disable or enable the port.

IPFC over FCR allows the routing of IPFC traffic between different EX_Ports where EX_Ports are connected to an edge fabric that has IPFC-capable devices. IPFC traffic is routed between EX and VEX Ports. IPFC traffic can be related to the same edge fabric or to different edge fabrics.

When the port is not active, the preferred domain ID is configurable. The preferred domain ID is used by the EX_Port's front phantom domain to request a domain ID from the principal switch. The domain ID received becomes the subsequent preferred domain ID, which is persistent and is displayed.

FCR supports Network OS edge (NOS) fabric connectivity when the EX_Port is set to Brocade Network OS fabric mode. No Integrated Routing License is required to support Network OS fabric connectivity, and selective device sharing will be allowed even when a license is not present.

This command is also used to configure the Inter-Chassis Link (ICL) ports to be EX_Ports. The ICL EX_Port is supported only on the VF-enabled 16 Gbps-capable chassis. Setting ICL EX_Port configuration in a non-VF switch or non-base switch will display an error.

All EX_Ports within a quad small form-factor pluggable (QSFP) must be present in the base switch and in disabled state for configuring ICL EX_Ports. If an EX_Port within the same QSFP is moved from the base switch to a logical switch, EX_Port configuration cannot be applied to any of the ports in the QSFP. When specifying an ICL port, all four E_Ports of the QSFP will be configured with the same EX_Port parameters.

Encryption and compression are not supported on an ICL EX_Ports. When this command is executed on ICL EX_Port without optional parameters, encryption and compression states are not displayed.

Notes The fabric ID must be the same for every router port connected to the same edge fabric, and different for every edge fabric. If two ports are connected to the same fabric but have been assigned different fabric IDs, one of them will be disabled due to a fabric ID oversubscription. If two fabrics have been assigned the same fabric ID, one of them will be disabled due to a fabric ID conflict.

When a port is changed from FL_Port to EX_Port, the topology is implicitly changed to point-to-point.

The front domain WWN field displays the WWN of the front domain. If the port is enabled and the state is "OK", the edge fabric principal switch domain ID and WWN also are displayed.

If the Fabric Parameter value is "Auto Negotiate", the port ID format, R_A_TOV, and E_D_TOV values display the negotiated values indicated by "(N)" next to them. The negotiated values are what the edge switch specifies in the ELP request. If the state is "Not OK", the R_A_TOV and E_D_TOV display "Not Applicable". By default, all EX_Ports are auto-ELP enabled.

If the Fabric Parameter attribute value is "User configured", the port ID format R_A_TOV and E_D_TOV values display the configured values.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands When invoked without operands, this command displays the usage. The following operands are supported:

- slot* On bladed systems only, specifies the slot number followed by a slash (/).
- port[-port]* Specifies a port or a port range, relative to the slot number on bladed systems, for example, 5/17-29. Use **switchShow** for a list of valid ports. When executed with [*slot*]*port* [*-port*] only, the command displays the current port configuration.
- a admin* Enables or disables the specified port as an EX_Port. Valid values are 1 (enable as EX_Port), 2 (disable as EX_Port and enable as non-EX_Port). **portCfgDefault** may also be used to disable EX_Ports.
- f fabricid* Specifies the fabric ID. The valid values for FID are from 1 through 128 or the alias name of the fabric. The alias name is displayed, only if alias name exists for the corresponding Edge Fabric ID. Execute **fcrconfigure --show** to view the alias name configuration.
- r ratov* Specifies the R_A_TOV used for port negotiation. Valid values are 2000 to 120000. This operand is only applicable if the "Fabric Parameter" attribute value is not "Auto Negotiate".
- e edtov* Specifies the E_D_TOV used for port negotiation. Valid values are 1000 to 60000. This operand is only applicable if the "Fabric Parameter" attribute value is not "Auto Negotiate".

- d *domainid*** Specifies the preferred domain ID. For Brocade native mode (**-m 0**), valid values are 1 to 239.
- p *pidformat*** Specifies the Port ID format. Valid values are 0-native, 1-core, 2-extended edge. This operand is applicable only when port mode is set to 0 (native mode). If port mode is not "Brocade Native", the Port ID format displays as "Not applicable".
- t *fabric_parameter*** Enables or disables negotiation of the fabric parameters. Valid values are 1 for enable and 2 for disable.
- m *port mode*** Specifies the port mode. The **-m** option enforces the same port mode for all the ports connected to the same edge fabric. If this option is selected, the port mode is compared against the online ports. If the modes are different, an error message is posted, and the command fails. Valid values are as follows:
- 0** Brocade Native mode.
- 5** Brocade Network OS fabric mode.
- i *mode*** Enables or disables Insistent Domain Id (IDID) for the specified EX_Port. Specify 1 to enable IDID, specify 2 to disable IDID. This command must be issued from a McDATA edge switch attached to a single or dual FCR configuration.
- help** Displays the command usage.

Examples

To set the fabric ID of port 2/1 to 5 and the port ID format to core:

```
switch:admin> portcfgexport 2/1 -f 5 -p 1
```

To configure port 2/0 to be an EX_Port and set the fabric ID to 4:

```
switch:admin> portcfgexport 2/0 -a 1 -f 4
```

To disable fabric parameter negotiation on port 2/0 of an EX_Port:

```
switch:admin> portcfgexport 2/0 -t 2
```

To enable IDID on port 2:

```
switch:admin> portcfgexport 2 -i 1
```

To enable Brocade Network OS (NOS) fabric mode on port 1/5:

```
switch:admin> portcfgexport 1/5 -a 1 -m 5
```

To enable Brocade Network OS (NOS) fabric mode using alias name:

```
switch:admin> portcfgexport 1/5 -a 1 -f Red_fabric
```

To view the configuration of an EX_Port:

```
switch:admin> portcfgexport 1/5
Port          1/5    info
Admin:                enabled
State:                OK
Pid format:          core(N)
Operate mode:        Brocade Native
Edge Fabric ID:      128
Alias name:          Red_fabric
Preferred Domain ID: 160
Front WWN:           50:00:51:e4:44:40:0e:80
Fabric Parameters:   Auto Negotiate
R_A_TOV:             10000(N)
E_D_TOV:             2000(N)
Authentication Type: DHCHAP
DH Group: 4
Hash Algorithm: SHA-1
Encryption: ON
Compression: ON
Forward Error Correction: ON
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

To view the configuration of an ICL EX_Port:

```
switch:admin> portcfgexport 5/12
Port          5/12    info
Admin:                enabled
State:                OK
Pid format:          core(N)
Operate mode:        Brocade Native
Edge Fabric ID:      11
Alias name:          orange_fabric
Front Domain ID:     160
Front WWN:           50:00:51:e4:8f:80:2e:0b
Principal Switch:    1
Principal WWN:       10:00:00:05:1e:48:f8:03
Fabric Parameters:   Auto Negotiate
R_A_TOV:             10000(N)
E_D_TOV:             2000(N)
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Forward Error Correction: ON
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

SEE ALSO [portCfgVEXPort](#), [portDisable](#), [portEnable](#), [portShow](#), [portCfgDefault](#), [fcrBcastConfig](#), [fcrConfigure](#)

portCfgFaultDelay

Configures the fault delay for a single FC port.

Synopsis `portcfgfaultdelay [slot/]port delay`

`portcfgfaultdelay --help`

Description Use this command to configure the fault delay of an FC port.

In the event that the link is noisy after a host power cycle, the switch may go into a soft fault state, which means a delay of R_A_TOV. Setting the mode value to 1 reduces the fault delay value to 1.2 seconds. The configuration is stored in nonvolatile memory and is persistent across switch reboots and power cycles.

Use the **portCfgShow** command to display the user-configured fault delay settings.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is applicable only to Fibre Channel ports.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

port

Specifies the number of the port to be configured, relative to its slot for bladed systems. Use **switchShow** for a listing of valid ports.

delay

Specifies the fault delay value for the port number. This operand is required. Valid values are one of the following:

0

Sets the fault delay to R_A_TOV (default).

1

Sets the fault delay to 1.2 seconds.

---help

Displays the command usage.

Examples To set the fault delay of a port to 1.2 seconds:

```
switch:admin> portcfgfaultdelay 2/3 1
```

To display the configuration

```
switch:admin> portcfgshow 2/3

Area Number:                21
Octet Speed Combo:          1(16G|8G|4G|2G)
Speed Level:                 AUTO(SW)
AL_PA Offset 13:           OFF
Trunk Port                   ON
Long Distance                OFF
VC Link Init                 OFF
Locked L_Port                OFF
Locked G_Port                OFF
Disabled E_Port              OFF
Locked E_Port                OFF
ISL R_RDY Mode               OFF
RSCN Suppressed              OFF
Persistent Disable           OFF
LOS TOV enable               OFF
NPIV capability              ON
QOS E_Port                   AE
Port Auto Disable:          OFF
Rate Limit                   OFF
EX Port                       OFF
Mirror Port                  OFF
Credit Recovery              ON
F_Port Buffers               OFF
Fault Delay:                 1(1.2sec)
NPIV PP Limit:               126
CSCTL mode:                  OFF
D-Port mode:                 OFF
D-Port over DWDM             OFF
Compression:                 OFF
Encryption:                   OFF
FEC:                          Active
```

```
switch:admin> portcfgshow
```

| Ports of Slot | 2 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|----------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Speed | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN | AN |
| Fill Word(On Active) | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| Fill Word(Current) | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| AL_PA Offset 13 | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Trunk Port | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON |
| Long Distance | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| VC Link Init | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Locked L_Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Locked G_Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Disabled E_Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| ISL R_RDY Mode | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| RSCN Suppressed | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Persistent Disable | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| LOS TOV enable | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| NPIV capability | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON |
| QOS E_Port | AE | AE | AE | AE | AE | AE | AE | AE | AE | AE | AE | AE | AE | .. | .. | .. | .. |
| EX Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Mirror Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Rate Limit | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Credit Recovery | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON | ON |

```

Fport Buffers      .. .. .
Port Auto Disable. .. .. .
CSCTL mode         .. .. .
D-Port mode        .. .. .
D-Port over DWDM   .. .. .
Compression        .. .. .
Encryption         .. .. .
FEC                .. .. .
Fault Delay        0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0
                    where AE:QoSAutoEnable, AN:AutoNegotiate, ..:OFF,
                    -:NotApplicable, ??:INVALID
(output truncated)

```

SEE ALSO [portCfgShow](#)

portCfgFec

Enables or disables Forward Error Correction (FEC) or FEC through Transmitter Training Signal (TTS) on 16 Gbps and 10 Gbps Fibre Channel links.

Synopsis

```
portcfgfec --enable [-FEC] [-TTS] [-force | -f] [slot/port [-port]]
portcfgfec --disable [-FEC] [-TTS] [-force | -f] [slot/port [-port]]
portcfgfec --show [slot/port [-port]]
portcfgfec --help
```

Description Use this command to enable or disable Forward Error Correction (FEC) or Transmitter Training Signal (TTS) on a specified port or on a range of ports, or to display the configuration.

FEC provides a mechanism for reducing error rates during data transmissions over 16 Gbps Fibre Channel links. When FEC is enabled on a port, the sender adds systematically generated error-correcting code (ECC) to its data transmission. This mechanism allows the receiver to detect and correct errors without needing to get additional information from the sender.

If the TTS mode is enabled, the port negotiates FEC through TTS. TTS and 64B/66B are mutually exclusive. At fixed speed, TTS is not compatible with 64B/66B at the peer port and the final port state will be "No_Sync". In auto negotiation, the local port will try the TTS encoding; if the other end is not compatible, then the local port will try the 8B/10B encoding and the port will settle in "N8" port speed if the peer port also supports 8B/10B encoding. Thus, the TTS mode should only be enabled if a similarly TTS-capable and enabled device is connected to the port.

The Brocade implementation of FEC is supported on 16 Gbps platforms and enables the switch to recover bit errors in 16 Gbps and 10 Gbps data streams. The FEC encoding can correct one burst of up to 11 error bits in every 2,112-bit transmission. The error correction covers both frames and primitives. There is no loss of bandwidth or added transmission data rate overhead to the 16 Gbps FC link.

By default, FEC without TTS is enabled switch-wide on all 16 Gbps platforms but FEC via TTS is not enabled. If FEC is already enabled on the ports, enabling FEC has no effect. If a range of ports is specified, some of which are already in the requested configuration, a notification is generated, and no action is taken for those ports only. All other ports in the specified range are updated. Enabling or disabling FEC is disruptive to traffic. The command prompts for confirmation to continue with enabling or disabling TTS or FEC via TTS for the online ports. You can override the prompt for confirmation with the **-force** option.

When used with the **--show** option, the command displays the following information for the specified ports:

| | |
|-----------------------|---|
| Port | The port index number |
| FEC Capable | Displays YES if the port supports FEC. Displays NO if the port does not support FEC. |
| FEC Configured | Displays ON if FEC is enabled on the port (default). Displays OFF if the feature is disabled. |

FEC via TTS Configured

Displays OFF if TTS is disabled on the port (default). Displays ON if the FEC negotiation via TTS feature is enabled.

FEC State

The FEC state can be active or inactive. An active FEC state indicates that FEC is enabled and actually running. An inactive state can indicate two conditions: FEC is enabled, but not running due to some error condition (for example, FEC may not be enabled on both ends of the link). Or FEC is disabled and therefore inactive.

Use the **portCfgShow** command to display the FEC configuration along with other port parameters. Use the **islShow** command to view interswitch link-level FEC configurations. Use the **portErrshow** and **portStatShow** commands to monitor data transmission errors. You should see a significant reduction in CRC errors on FEC-enabled links.

Notes FEC is supported on the following links:

- Between E_Ports on all 16 Gbps platforms running Fabric OS v7.0.0 or later. Both sides of the link must be configured with port speeds of 10 Gbps and 16 Gbps.
- Between F_Ports and N_Ports in Access Gateway mode (requires Fabric OS v7.1.0 and later on the AG and the switch).
- Between Brocade 16 Gbps capable HBAs (Catapult2) Host Bus Adapters and an F_Port. The HBA must be running v3.2 or later and the switch must be running Fabric OS v7.1.0.

FEC is compatible with QoS, Credit Recovery, and Fabric-Assigned Port WWM (FA-PWWN).

FEC is not supported on D_Ports configured with Dense Wavelength Division Multiplexing (DWDM).

Starting from Fabric OS v7.3.0, the TTS mode is supported only for F_Ports. If a port initializes as an E_Port, it is disabled with a warning message and its peer port will be in "No_Light". status.

The TTS mode is not supported for ICL ports, EX_Ports, D_Ports, SIM ports, Mirror ports, Software ASN, fixed 1 Gbps, fixed 2 Gbps, fixed 4 Gbps, or fixed 8 Gbps ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

On bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port[-port]

Specifies a port or a port range, relative to the slot number on bladed systems, for example, 5/17-29. Multiple port ranges are not supported with this command.

--enable [-FEC] [-TTS]

Enables FEC, FEC through TTS or both on the specified ports. Use the **-FEC** option to enable FEC only, use the **-TTS** option to enable TTS only, or use both **-FEC** and **-TTS** options together to enable FEC through TTS.

--disable [-FEC] [-TTS]

Disables FEC, FEC through TTS on the specified ports. Use the **-FEC** option to disable FEC only, use the **-TTS** option to disable TTS only, or use both **-FEC** and **-TTS** options together to disable FEC through TTS.

-force | -f

Enables or disables FEC or FEC through TTS without prompting for a confirmation.

--show

Displays the FEC and TTS configurations on the specified ports.

--help

Displays the command usage.

Examples

To enable FEC on a single port and to display the configuration:

```
switch:admin> portcfgfec --enable -FEC 5/28
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
switch:admin> portcfgfec --show 5/28
Port: 412
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: OFF
FEC State: Active
```

To enable FEC on a port:

```
switch:admin> portcfgfec --enable -FEC 18
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
```

To enable FEC on a port range:

```
switch:admin> portcfgfec --enable -FEC 0-8
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been enabled.
```

To enable the FEC feature on a range of ports, some of which were previously enabled (the following example enables port 4):

```
switch:admin> portcfgfec --enable -FEC 2-4
Warning : FEC changes will be disruptive to the traffic
Same configuration for port 2
Warning : FEC changes will be disruptive to the traffic
Same configuration for port 3
Warning : FEC changes will be disruptive to the traffic
Same configuration for port 4
```

To disable the FEC feature on a port range:

```
switch:admin> portcfgfec --disable -FEC 0-8
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
Warning : FEC changes will be disruptive to the traffic
FEC has been disabled.
```

To enable TTS on a port:

```
switch:admin> portcfgfec --enable -TTS 7
WARNING: 16G FEC with TTS is only supported on F-Ports. \
        Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
TTS has been enabled.
```

To display the TTS on a port:

```
switch:admin> portcfgfec --show 7
Port: 7
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: ON
FEC State: Active
```

To disable TTS on a port:

```
switch:admin> portcfgfec --disable -TTS 1
WARNING: 16G FEC with TTS is only supported on F-Ports. \
        Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
TTS has been disabled.
```

To enable both FEC and TTS on a port:

```
switch:admin> portcfgfec --enable -FEC -TTS 2
WARNING: 16G FEC with TTS is only supported on F-Ports. \
    Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
FEC & TTS have been enabled for port 2.
```

To disable both FEC and TTS on a port:

```
switch:admin> portcfgfec --disable -FEC -TTS 2
WARNING: 16G FEC with TTS is only supported on F-Ports. \
    Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
FEC & TTS have been disabled for port 2.
```

To enable TTS on a port range using the force option:

```
switch:admin> portcfgfec --enable -TTS 1-5 -f
FEC & TTS have been enabled for port 1.
FEC & TTS have been enabled for port 2.
FEC & TTS have been enabled for port 3.
FEC & TTS have been enabled for port 4.
FEC & TTS have been enabled for port 5.
```

To enable TTS feature on a range of ports, which were previously enabled using the force option:

```
switch:admin> portcfgfec --enable -TTS 16-21 -f
Same configuration for port 16
Same configuration for port 17
Same configuration for port 18
TTS has been enabled.
TTS has been enabled.
TTS has been enabled.
```

To enable FEC and TTS feature when FEC is previously enabled:

```
switch:admin> portcfgfec --enable -FEC -TTS 1
WARNING: 16G FEC with TTS is only supported on F-Ports. \
    Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
Same FEC configuration; and TTS has been enabled for port 1.
```

To enable FEC and TTS feature when TTS is previously enabled:

```
switch:admin> portcfgfec --enable -FEC -TTS 1
WARNING: 16G FEC with TTS is only supported on F-Ports. \
    Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
Same TTS configuration; and FEC has been enabled for port 1.
```


To disable FEC and TTS feature when FEC is previously disabled and TTS is enabled:

```
switch:admin> portcfgfec --disable -FEC -TTS 1
WARNING: 16G FEC with TTS is only supported on F-Ports. \
        Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
Same FEC configuration; and TTS has been disabled for port 1.
```

To disable FEC and TTS feature when TTS is previously disabled and FEC is enabled:

```
switch:admin> portcfgfec --disable -FEC -TTS 1
WARNING: 16G FEC with TTS is only supported on F-Ports. \
        Other port types will be disabled.
FEC changes are disruptive. Are you sure you want to continue?
(yes, y, no, n): [no]: y
Same TTS configuration; and FEC has been disabled for port 1.
```

To display the FEC and TTS configuration on range of ports:

```
switch:admin> portcfgfec --show 7-10
Port: 7
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: ON
FEC State: Active

Port: 8
FEC Capable: YES
10G/16G FEC Configured: OFF
16G FEC via TTS Configured: OFF
FEC State: Inactive

Port: 9
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: OFF
FEC State: Inactive

Port: 10
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: OFF
FEC State: Inactive
```

SEE ALSO [islShow](#), [portCfgShow](#), [portErrShow](#), [portStatsShow](#)

portCfgFillword

Configures the fill word for a single 8G FC port.

Synopsis `portcfgfillword [slot/]port, mode [passive]`

`portcfgfillword --help`

Description Use this command to configure the fill word of an 8G FC port. By default, this command disables and re-enables the port and the port comes online with the new fill word setting. When passive option **1** is used after the mode, the new settings are not applied until the next time the port goes offline and comes back online. When passive option **0** is used after the mode, the new settings are applied immediately. The configuration is stored in nonvolatile memory and is persistent across switch reboots or power cycles.

Notes This command is applicable only to 8G FC ports on 16 Gbps-capable platforms.

This configuration cannot be set on VE_Ports or VEX_Ports.

Use the **portCfgShow** command to display user-configured fill word settings.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

port

Specifies the number of the port to be configured, relative to its slot for bladed systems. Use **switchShow** for a listing of valid ports.

mode

Specifies the fill word for the port number. This operand is required. Valid values are one of the following:

0 | -idle-idle

Sets IDLE mode in the Link Init and IDLE as the fill word (default).

1 | -arbff-arbff

Sets ARB(ff) in the Link Init and ARB(ff) as the fill word.

2 | -idlef-arbff

Sets IDLE mode in the Link Init and ARB(ff) as the fill word.

3 | -aa-then-ia

Attempts hardware arbff-arbff (mode 1) first. If the attempt fails to go into active state, this command executes software idle-arb (mode 2). Mode 3 is the preferable to modes 1 and 2 as it captures more cases.

passive

Specifies the fill word configuration changes to take effect immediately or next port toggle. This operand is optional. Valid values are one of the following:

0

Specify 0 for the new settings to take effect immediately.

1

Specify 1 to postpone the new settings to take effect only after the next time the port goes offline and comes back up. It prevents immediate application of the new settings through a disruptive **portDisable** and **portEnable** operation.

Examples

To set the fill word of a port to ARBFF-ARBFF using the numeric mode notation:

```
switch:admin> portcfgfillword 2/3, 1
```

To set the fill word of a port to ARBFF-ARBFF using the -arbff-arbff option:

```
switch:admin> portcfgfillword 2/3, -arbff-arbff
```

To set the fill word of a port to IDLE-ARBFF using the -idle-arbff option:

```
switch:admin> portcfgfillword 27, -idle-arbff
switch:admin> portcfgshow 27
Area Number:                27
Speed Level:                 AUTO(HW)
Fill Word(On Active)        2(SW Idle-Arbff)
Fill Word(Current)          2(SW Idle-Arbff)
(output truncated)
```

To set the fill word of a port using the passive option:

```
switch:admin> portcfgfillword 10/2 2 1

switch:admin> portcfgshow 10/2
Area Number:                82
Speed Level:                 AUTO(HW)
Fill Word(On Active)        2(SW Idle-Arbff)
Fill Word(Current)          1(Arbff-Arbff)
(output truncated)

switch:admin> portdisable 10/2; portenable 10/2
switch:admin> portcfgshow 10/2
Area Number:                82
Speed Level:                 AUTO(HW)
Fill Word(On Active)        2(SW Idle-Arbff)
Fill Word(Current)          2(SW Idle-Arbff)
AL_PA Offset 13:            OFF

switch:admin> portcfgfillword 10/2 3 0
switch:admin> portcfgshow 10/2
Area Number:                82
Speed Level:                 AUTO(HW)
Fill Word(On Active)        3(A-A then SW I-A)
```

portCfgFillword

| | |
|--------------------|--------------------|
| Fill Word(Current) | 3(A-A then SW I-A) |
| AL_PA Offset 13: | OFF |
| Trunk Port | OFF |
| Long Distance | OFF |

SEE ALSO [portCfgShow](#)

portCfgFlogiLogout

Enables the Base Device Logout functionality on the port.

Synopsis `portcfgflogilogout --enable [slot]port | -all`

`portcfgflogilogout --disable [slot]port | -all`

`portcfgflogilogout --help`

Description Use this command to enable or disable Base Device Logout functionality on a specified port or port range. By default, the functionality is disabled on all the ports.

This functionality allows NPIV devices to remain logged in even after the base device logs out.

A base device is a device on a F_Port which has the base PID. The base device logs in with a FLOGI.

The execution of this command is disruptive. The online ports will toggle with this command but prompts for user confirmation if any port is online.

You cannot configure the Base Device Logout functionality if the standby device is running on a lower version of the firmware.

The feature is not supported on ICL ports. You cannot enable this feature on the ports that do not support NPIV capability or if NPIV is disabled by configuration.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

`[slot]port | -all`

Specifies a port number. You can specify a port or range of ports. You can use `-all` to include all ports on the logical switch.

`--enable`

Enables Base Device Logout functionality on the specified port.

`--disable`

Disables Base Device Logout functionality on the specified port.

`--help`

Displays the command usage.

Examples To enable Base Device Logout functionality on a port:

```
switch:admin> portcfgflogilogout --enable 2/1
```

To disable Base Device Logout functionality on a port:

```
switch:admin> portcfgflogilogout --disable 2/1
```

SEE ALSO [portCfgNPIVPort](#)

portCfgFportBuffers

Configures F_Port buffer allocation.

Synopsis `portcfgfportbuffers --enable [slot]port buffers`

`portcfgfportbuffers --disable [slot]port`

Description Use this command to change the default buffer allocation for an F_Port and to allocate a specified number of buffers to the port. When port buffer allocation is enabled, the number of buffers specified override the default F_Port buffer allocation. When the configuration is disabled, the default buffer allocation is restored. Only an F_Port can utilize the buffers allocated by this command, and the allocated buffers are reserved only for this port.

The F_Port buffer configuration is persistent across system reboots.

Use the **portBufferShow** command to determine current port buffer allocations.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The F_Port buffer feature does not support ports configured as EX_Ports, Mirror Ports, Long Distance Ports, L_Ports, QoS Ports, Fast Write, and Trunk Areas.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands This command has the following operands:

--enable

Enables the F_Port buffer configuration on a specified port. A port and buffer allocation must be specified with this option.

--disable

Disables the F_Port buffer configuration on a specified port.

slot

Specifies the slot number on bladed systems, followed by a slash (/).

buffers

Specifies the number of buffers to be allocated to the specified port. The specified buffer allocation takes effect when the F_Port comes online. This operand is required with the **--enable** option. The minimum buffer allocation is the default number of buffers plus 1. The maximum is determined by the remaining buffer allocations in the port's port group. Use **portBufferShow** to determine the number of remaining free buffers.

Examples To allocate 12 buffers to an F_Port:

```
switch:admin> portcfgfportbuffers --enable 2/44 12
```

To disable the port buffer configuration and return to the default buffer allocation:

```
switch:admin> portcfgfportbuffers --disable 2/44 12
```

SEE ALSO [portBufferShow](#)

portCfgGE

Manages the port configuration of the GbE/10GbE ports.

Synopsis

```
portcfgge [slot/] port --set -speed speed
portcfgge [slot/] port --set -mediatype media
portcfgge [slot/] port --set -channel channel_num
portcfgge [slot/] port --set -lan
portcfgge [slot/] port --set -wan
portcfgge [slot/] port --enable -autoneg
portcfgge [slot/] port --disable -autoneg
portcfgge [slot/][ port] --show [-lmac]
portcfgge --help
```

Description Use this command to manage the port configuration of the GbE/10GbE ports. This command switches the port speed between 1G and 10G or switch the auto-negotiate mode. The auto-negotiate option allows to control auto-negotiate on the PHY for 1G mode.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|---------------------|--|
| <i>slot</i> | For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/). |
| <i>port</i> | Specifies the number of the GbE port to be configured. On the Brocade FX8-24 blade, the GbE ports are numbered ge0 - ge9 and 10 GbE ports are numbered xge0 - xge1. On the Brocade 7840 switch and the Brocade SX6 extension blade, 40 GbE ports are numbered ge0 - ge1 and 1G/10GbE ports are numbered ge2 - ge17. Use the switchShow command for a list of valid ports. |
| --set | Configures the GE ports. |
| -speed speed | Configures the port speed for the GE ports. The valid speeds are 1G or 10G. |

- mediatype *media***
Specifies transmit mode for the GE ports. The valid types are optical or copper. Specifying optical selects the optical mode for the GbE port. Specifying copper selects a copper mode for the GbE port.
- channel *channel_num***
Configures the tunable small form-factor pluggable (TSFP) channel ID for the 10 GE ports. The valid channel ID range is 1 through 102. This option is supported only on the Brocade 7840 switch and the Brocade SX6 extension blade.
- lan**
Configures a GE port as a LAN port. This option is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- wan**
Configures a GE port as a WAN port. This option is supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- enable -autoneg**
Enables the auto negotiation mode for the GE port.
- disable -autoneg**
Disables the auto negotiation mode for the GE port.
- show**
Displays the current GE port configurations.
- lmac**
Displays the local MAC address. This operand is optional.
- help**
Displays the command usage.

Examples

To display the current GE port configurations:

```
switch:admin> portcfgge --show
Port          Speed   Flags   LAG-ID
-----
ge0           40G    A----   -
ge1           40G    A----   -
ge2           10G    A-LG-   lag0
ge3           10G    A-LG-   lag0
ge4           10G    A----   -
ge5           10G    A----   -
ge6           10G    A----   -
ge7           1G     --LG-   edgeSw1
ge8           10G    A----   -
ge9           10G    A----   -
ge10          10G    A----   -
ge11          10G    A----   -
ge12          10G    A----   -
ge13          10G    A----   -
ge14          10G    A----   -
ge15          10G    A----   -
ge16          10G    A----   -
ge17          10G    A----   -
-----
```


To configure the port speed to 1G for ge8 and verify the configuration:

```
switch:admin> portcfgge ge8 --set -speed 1G
switch:admin> portcfgge --show
Port          Speed   Flags   LAG-ID
-----
ge0           40G    A----  -
ge1           40G    A----  -
ge2           10G    A-LG-  lag0
ge3           10G    A-LG-  lag0
ge4           10G    A----  -
ge5           10G    A----  -
ge6           10G    A----  -
ge7           1G     --LG-  edgeSw1
ge8           1G     A----  -
ge9           10G    A----  -
ge10          10G    A----  -
ge11          10G    A----  -
ge12          10G    A----  -
ge13          10G    A----  -
ge14          10G    A----  -
ge15          10G    A----  -
ge16          10G    A----  -
ge17          10G    A----  -
-----
```

To configure a GE port for LAN operation and verify the configuration:

```
switch:admin> portcfgge ge10 --set -lan
Operation Succeeded.
SB118:FID128:root> portcfgge --show
Port          Speed   Flags   LAG-ID
-----
ge0           40G    A----  -
ge1           40G    A----  -
ge2           10G    A-LG-  lag0
ge3           10G    A-LG-  lag0
ge4           10G    A----  -
ge5           10G    A----  -
ge6           10G    A----  -
ge7           1G     --LG-  edgeSw1
ge8           10G    A----  -
ge9           10G    A----  -
ge10          10G    A-L--  -
ge11          10G    A----  -
ge12          10G    A----  -
ge13          10G    A----  -
ge14          10G    A----  -
ge15          10G    A----  -
ge16          10G    A----  -
ge17          10G    A----  -
-----
Flags: A:Auto-Negotiation Enabled  C:Copper Media Type
      L:LAN Port  G=LAG Member
-----
```

SEE ALSO [portCfg](#), [portCfgDefault](#), [portShow](#)

portCfgGeMediaType

Sets the selected mode for the ge0 and ge1 ports on the Brocade 7840.

Synopsis `portcfgmediatype ge_port [optical | copper]`

`portcfgmediatype --help`

Description Use this command to set the mode for ge0 and ge1 ports on the Brocade 7840.

Each of the ge0 and ge1 ports on the Brocade 7840 can be configured as a copper port (RJ45) or an optical port (SFP). The copper ports (RJ45) can only accept copper connections. The optical ports can accept any Brocade-branded SFP, optical or copper. Each of the ge0 and ge1 ports can be configured either as a copper port or an optical port, but not both.

The following configurations are possible:

Default

- ge0 (RJ45) - active and ge0 (SFP) - inactive
- ge1 (RJ45) - active and ge1 (SFP) - inactive

One copper and one optical port active

- ge0 (RJ45) - active and ge0 (SFP) - inactive
- ge1 (RJ45) - inactive and ge1 (SFP) - inactive

Or:

- ge1 (RJ45) - inactive and ge1 (SFP) - active
- ge0 (RJ45) - active and ge0 (SFP) - active

Both copper ports inactive. Both optical SPP ports active

- Ge0 (RJ45) - inactive and Ge0 (SFP) - active
- Ge1 (RJ45) - inactive and Ge1 (SFP) - active

The copper-only GbE ports are selected by default. To select ge0 or ge1 as optical mode specify the operand as **optical**. Specify **copper** to re-select the default copper port.

When used with the *ge_port* operand only, this command displays the currently selected GbE port and media type.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported only on the Brocade 7840 and only on ge1 and ge0 ports. On all other platforms or ports, this command returns a "not supported message."

Operands This command has the following operands:

ge_Port

Specifies the GbE port to be selected. Valid ports are ge0 and ge1. This operand is required.

optical | copper

Specifies which mode to select. Specifying **optical** selects the optical mode for the GbE port. Specifying **copper** selects a copper mode for the GbE port. This operand is optional; if omitted, the current media type is displayed for the specified GbE port.

--help

Displays the command usage.

Examples

To select an optical mode for ge0:

```
switch:admin> portcfgmediatype ge0 optical
```

To display the current configuration:

```
switch:admin> portcfgmediatype ge0  
Port ge1 is configured in optical mode
```

To restore the default configuration (copper only ports are selected):

```
switch:admin> portcfgmediatype ge0 copper  
switch:admin> portcfgmediatype ge0  
Port ge0 is configured in copper mode
```

SEE ALSO [portCfgShow](#)

portCfgGport

Designates a port as a G_Port; removes G_Port designation.

Synopsis `portcfggport [slot]/port,mode`

Description Use this command to designate a port as a G_Port. After successful execution of this command, the switch attempts to initialize the specified port as an F_Port only, and does not attempt loop initialization (FL_Port) on the port. A port designated as a G_Port can become an E_Port. This configuration can be cleared but not set on VE/VEX_Ports. Changes made by this command are persistent across switch reboots or power cycles.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the port to be configured, followed by a slash (/).

port

Specify the port to be configured, relative to its slot for bladed systems. Use **switchShow** to display a list of valid ports.

mode

Specify a value of 1 to designate the port as a G_Port or specify a value of 0 to remove the G_Port designation from the port. A value of 0 is the default port state. Mode must be preceded by a comma. This operand is required.

Examples To configure port as a locked G_Port:

```
switch:admin> portcfggport 1/3,1
```

SEE ALSO [configure](#), [portCfgLport](#), [portShow](#), [switchShow](#)

portCfgISLMode

Enables or disables ISL R_RDY mode on a port.

Synopsis `portcfgislmode [slot/]port,mode`

Description Use this command to enable or disable interswitch link read-ready (ISL R_RDY) mode on a port. Use the **portCfgShow** command to determine whether ISL R_RDY mode is enabled on a port.

In ISL R_RDY mode, the port sends a primitive signal that the port is ready to receive frames. The port sends an exchange link parameter (ELP) with flow control mode O2. If a port is ISL R_RDY enabled, it can only receive an ELP with flow control mode O2. A received ELP with flow control mode O1 will segment the fabric.

This mode cannot detect any inconsistencies in fabric operating mode parameters, such as the PID format of connected ports. Before enabling ISL R_RDY mode, ensure that all fabric-wide parameters are consistent for every switch in the fabric.

Use **configShow fabric.ops** to view a complete listing of fabric operating mode parameters on the switch.

The following E_Port configurations are not applicable to a port configured for ISL R_RDY mode. If configured, these port configuration parameters are ignored during E_Port initialization:

- Trunk port
- VC link init

The **portCfgISLMode** and **portCfgLongDistance** levels LE, LD, or LS only can be enabled at the same time. Such an ISL uses R_RDY mode of flow control over the long distance link. This feature is not backward compatible with firmware versions that do not support it.

Notes Changes made by **portCfgISLMode** are persistent across switch reboots and power cycles.

This configuration can be cleared but not set on VE/VEX_Ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands This command has the following operands:

| | |
|-------------|--|
| <i>slot</i> | For bladed systems only, specify the slot number of the port to display, followed by a slash (/). |
| <i>port</i> | Specify the port to display, relative to its slot for bladed systems. Use switchShow to list valid ports. |
| <i>mode</i> | Specify 1 to enable ISL R_RDY mode. Specify 0 to disable ISL R_RDY mode. |

portCfgISLMode

Examples To enable ISL R_RDY mode on a port:

```
switch:admin> portcfgislmode 1/3, 1
ISL R_RDY Mode is enabled for port 3. Please make sure
the PID formats are consistent across the entire fabric.
```

To disable ISL R_RDY mode on a port:

```
switch:admin> portcfgislmode 1/3, 0
```

SEE ALSO [configure](#), [portCfgLongDistance](#), [portCfgShow](#)

portCfgLongDistance

Configures a port to support long distance links.

Synopsis `portcfglongdistance [slot/]port
[distance_level]
[vc_translation_link_init]
[-distance distance] | [-buffer buffers]
[-framesize frame_size]
[-fecenable | -fecdisable]`

Description Use this command to allocate frame buffer credits to a port or to configure a specified long distance link. The port can only be used as an E_Port. Changes made by this command are persistent across switch reboots and power cycles. This configuration can be cleared but not set on VE/VEX_Ports.

Long distance configuration allows native FC ports to run WAN/LAN connections. It ensures that the full bandwidth of a link or trunk can be utilized for a particular long distance configuration. The receiving port must have sufficient buffers available, so that the transmitting port can stuff the link with enough frames to fill the entire length of the link. As the distance between switches and the link speed increases, additional buffer-to-buffer credits are required to maintain maximum performance. If a port is configured as a long distance port, the remaining ports of that port group could be disabled, fail to initialize, or move to "buffer limited" mode due to a lack of frame buffer credits.

The number of credits reserved for a port depends on the switch model and on the extended fabric mode for which it is configured. Not all distance modes are supported by all platforms. Refer to the *Fabric OS Administrator's Guide* for details on platform-specific buffer credit models, long distance mode support, and maximum distance supported for specific hardware configurations.

Notes This command requires an Extended Fabrics license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

A long-distance link can also be configured to be part of a trunk group. Refer to **portCfgTrunkPort** help for details.

When a port is configured as a long-distance port, the output of **portShow** and **switchShow** displays the long-distance level. Refer to **portShow** help and **switchShow** help for details.

The **portCfgISLMode** and **portCfgLongDistance** LE, LD, or LS levels can be enabled at the same time. Such an ISL uses the R_RDY mode of flow control over the long distance link. While using R_RDY mode flow control, an E_Port cannot form trunk groups of long-distance links even if the trunking is enabled. This feature is not backward compatible with firmware versions that do not support it.

Ctrl-D cancels the configuration update.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands This command has the following operands:

slot Specifies the slot number (for bladed systems only), followed by a slash (/).

portCfgLongDistance

port

Specifies the number of the port to be configured relative to its slot for bladed systems. Use **switchShow** to display a list of valid ports. This operand is required.

distance_level

Specifies the long distance level as one of the following (the numerical value representing each *distance_level* is shown in parentheses):

L0 (0)

Specifies L0 to configure the port as a regular port. A total of 20 full-size frame buffers are reserved for data traffic, regardless of the port's operating speed; therefore, the maximum supported link distance is up to 5 km at 2 Gbps, up to 2 km at 4 Gbps and up to 1 km at 8, 10, and 16 Gbps.

LE (3)

Specifies LE mode to configure an E_Ports distance greater than 5 km and up to 10 km. A total of 5, 10, 20, 40, 50, and 80 full-size frame buffers are reserved for data traffic at port speeds of 1 Gbps, 2 Gbps, 4 Gbps, 8, 10, and 16 Gbps. LE does not require an Extended Fabrics license. If a port in LE mode is set to autonegotiation, it will reserve the buffers for the highest support speed on that port. If this is not the desired buffer allocation, you should set the port to a fixed speed.

LD (5)

Specifies LD for automatic long-distance configuration. The buffer credits for the given E_Port are automatically configured based on the actual link distance. Up to a total of 1452 full-size frame buffers are reserved depending on the distance measured during E_Port initialization. The *desired_distance* is used as the upper limit to the measured distance. When the *desired_distance* is less than the measured distance, the port will come up in "Buffer Limited" mode and the port takes the buffers calculated from the desired distance. If a port in LD mode is set to autonegotiation, it will reserve the buffers for the highest support speed on that port. If this is not the desired buffer allocation, you should set the port to a fixed speed.

LS (6)

Specifies LS mode to configure a static long distance link with a fixed buffer allocation greater than 10 km. Up to a total of 1452 full-size frame buffers are reserved for data traffic, depending on the specified *desired_distance* value. If a port in LS mode is set to autonegotiation, it will reserve the buffers for the highest support speed on that port. If this is not the desired buffer allocation, you should set the port to a fixed speed.

vc_translation_link_init

Specifies the fill words used on long distance links. When set to 1, the link uses ARB fill words (default). When set to 0, the link uses IDLE fill words. The IDLE fill word option is not compatible with QoS configured links and Credit Recovery enabled links. You must disable these features before configuring long distance IDLE fill words.

-distance *distance*

This parameter is required when a port is configured as an LD or an LS mode link. In LD mode, the value of `desired_distance` is the upper limit of the link distance and is used to calculate buffer availability for other ports in the same port group. When the measured distance exceeds the value of *distance*, this value is used to allocate the buffers. In this case, the port operates in degraded mode instead of being disabled due to insufficient buffers. In LS mode, the actual link distance is not measured, instead the *distance* is used to allocate the buffers required for the port.

-framesize *frame_size*

Specifies the average frame size for LD and LS Mode.

-buffer *buffers*

Specifies the desired buffer for LD and LS mode.

-fcenable

Enables Forward Error Correction on 10 Gbps and 16 Gbps speed ports.

-fecdisable

Disables Forward Error Correction on 10 Gbps and 16 Gbps speed ports.

FEC via TTS is always enabled on 32 Gbps speed ports.

Examples

To configure a switch port 118 to support a 100 km link and be initialized using the long distance link initialization protocol:

```
switch:admin> portcfglongdistance 12/6 LS 1 -distance 100
switch:admin> portshow 12/6
portIndex: 118
portName: slot12 port6
portHealth: No Fabric Watch License

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x103 PRESENT ACTIVE E_PORT G_PORT U_PORT \
  SEGMENTED LOGIN
LocalSwcFlags: 0x0
portType: 24.0
portState: 1 Online
Protocol: FC
portPhys: 6 In_Sync portScn: 64 Segmented Flow control \
  mode 0
port generation number: 14
state transition count: 12

portId: 017600
portIfId: 43c2001e
portWwn: 20:76:00:05:1e:e5:cb:00
portWwn of device(s) connected:

Distance: static (desired distance = 100 Km)
portSpeed: N8Gbps

FEC: Inactive
Credit Recovery: Inactive
LE domain: 0
```

portCfgLongDistance

```
FC Fastwrite: OFF
Interrupts:    0    Link_failure: 0    Frjt: 0
Unknown:      0    Loss_of_sync: 0    Fbsy: 0
Lli:          70    Loss_of_sig: 0
Proc_rqrd:   205    Protocol_err: 0
Timed_out:    0    Invalid_word: 0
Rx_flushed:   0    Invalid_crc: 0
Tx_unavail:   0    Delim_err: 0
Free_buffer:  0    Address_err: 0
Overrun:      0    Lr_in: 9
Suspended:    0    Lr_out: 5
Parity_err:   0    Ols_in: 0
2_parity_err: 0    Ols_out: 6
CMI_bus_err:  0
```

To configure desired buffers:

```
switch:admin> portcfglongdistance 2/35 \  
                LS 1 -buffers 400  
Reserved Buffers = 420
```

To configure average frame size:

```
switch:admin> portcfglongdistance 2/35 \  
                LS 1 -distance 100 -framesize 1024
```

SEE ALSO [configure](#), [portCfgISLMode](#), [portCfgTrunkPort](#), [portCfgShow](#), [portShow](#), [switchShow](#)

portCfgLossTov

Enables or disables de-bouncing of signal loss for front end ports.

Synopsis

```
portcfglosstov [slot]/port [-disable | 0]
portcfglosstov [slot]/port [-enable_fixed | 1]
portcfglosstov [slot]/port [-enable_all | 2]
portcfglosstov -help
```

Description Use this command to enable or disable the de-bouncing of loss of signal for 100 ms for front end ports. Use **portCfgShow** to display the current setting.

If executed without operands, the command prints the usage.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

On a chassis, this command must be executed on the active CP.

Operands This command has the following operands:

| | |
|--------------------------|--|
| <i>slot</i> | For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/). |
| <i>port</i> | Specifies the port to be configured, relative to its slot for bladed systems. Use switchShow to list valid ports. |
| -disable 0 | Disables the configuration. |
| -enable_fixed 1 | Enables the configuration on fixed speed ports. |
| -enable_all 2 | Enables the configuration on autonegotiated and fixed speed ports. |
| -help | Displays the command usage. |

Examples To enable the configuration on fixed speed port 1/5 and to display the configuration:

```
switch:admin> portcfglosstov 1/5 1

switch:admin> portcfgshow 1/5
Area Number:          5
Speed Level:          AUTO (HW)
Fill Word(On Active): 0 (Idle-Idle)
Fill Word(Current):   0 (Idle-Idle)
AL_PA Offset 13:      OFF
```

portCfgLossTov

| | |
|--------------------|--------------|
| Trunk Port | ON |
| Long Distance | OFF |
| VC Link Init | OFF |
| Locked L_Port | OFF |
| Locked G_Port | OFF |
| Disabled E_Port | OFF |
| Locked E_Port | OFF |
| ISL R_RDY Mode | OFF |
| RSCN Suppressed | OFF |
| Persistent Disable | OFF |
| LOS TOV mode | 1(ON:FixSpd) |
| (Output truncated) | |

To disable the configuration on port 1/5:

```
switch:admin> portcfglossstov 1/5 0
```

SEE ALSO [portCfgShow](#)

portCfgLport

Configures a port as an L_Port.

Synopsis `portcfglport`
`[[slot/]port] locked_mode`
`[private_mode] [duplex_mode]`

Description Use this command to designate a port as an L_Port, and to configure its behavior. When a port is designated as an L_Port, the switch attempts to initialize that port as a fabric L_Port (FL_Port). The switch will never attempt a point-to-point (F_Port) initialization on the port. By default the L_Port will be a public L_Port. It can be configured as a private L_Port, in which case it will reject fabric login (FLOGI).

Notes This configuration can be cleared but not set on VE/VEX_Ports.

VE_Ports on the Brocade FX8-24 platforms do not support Arbitrated Loops. If one of the virtual FC ports (16-31) is preconfigured as an L_Port, and a Brocade FX8-24 Extension blade is inserted into the slot, the configuration is not honored. The system logs a RASLOG error and the port defaults back to a G_Port. Attempts to configure a VE_Port as an L_Port are rejected with a corresponding system message.

The Fabric OS port configuration commands are not supported on FCoE ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands When invoked without operands, this command reports the L_Port conditions for all ports present. The following operands are supported:

slot For bladed systems only, specify the slot number of the port to be configured, followed by a slash (/).

port Specify a port number to be configured, relative to its slot for bladed systems. Use **switchShow** to display a list of valid ports. This operand is optional; if omitted, this command displays the L_Port conditions for all ports.

locked_mode Specify 1 to configure the specified port as a locked L_Port. Specify 0 to remove the locked L_Port configuration from the port (default). This operand is required, if a port is specified.

private_mode Specify 1 to configure the L_Port as a private L_Port. Specify 0 to configure the L_Port as a public L_Port. This operand is optional; if omitted, the default value of 0 is used. On unsupported platforms, this command is blocked and returns an appropriate message.

duplex_mode Specify 2 to configure the specified port as a full-duplex L_Port with fairness. Specify 0 to configure the L_Port as a full-duplex L_Port. This operand is optional; if omitted, the default value of 0 is used.

portCfgLport

Examples To configure ports 8, 14, and 15 as locked L_Ports:

```
switch:admin> portcfglport 4/8 4/14-15, 1
```

To display the L_Port conditions:

```
switch:admin> portcfglport
[...]
```

| Ports of Slot | 4 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|----------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Locked L_Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | ON | .. | .. | .. | .. | .. | ON | ON |
| Private L_Port | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Locked Loop HD | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |
| Loop Fairness | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. | .. |

```
(output truncated)
```

SEE ALSO [configure](#), [portCfgShow](#), [portShow](#), [switchShow](#)

portCfgNonDfe

Enables or disables non-Decision Feedback Equalization (DFE) mode on 8 Gbps Fibre Channel links.

Synopsis

```
portcfgnondfe --enable [slot/]port[-port]
portcfgnondfe --disable [slot/]port[-port]
portcfgnondfe --show [slot/]port[-port]
portcfgnondfe --help
```

Description Use this command to enable or disable non-DFE mode on a specified port or on a range of ports, or to display the configuration and state.

Active receiver DFE uses sophisticated algorithms to automatically adjust the receiver to compensate for signal distortions.

If enabled, this command disables the active receiver DFE and instead uses a fixed receiver equalization. The enabled mode is further indicated by the suffix of "Static" in the **--show** option.

If the non-DFE enabled port is connected to a port that does not require fixed receiver equalization, CRC errors may be detected on the port and the link may toggle.

By default, non-DFE mode is disabled, which means that automatic receiver adjustment through DFE is activated. Even if disabled, non-DFE is still automatically activated if both of the following conditions are met:

- Port speed is 8G or N8; and
- The received fillword is IDLE.

The disabled mode is further indicated by the suffix of "Auto" in the **--show** option.

If non-DFE is already enabled on the ports, enabling non-DFE has no effect. If a range of ports is specified, some of which are already in the requested configuration, a notification is generated, and no action is taken for those ports only. All other ports in the specified range are updated. Enabling or disabling non-DFE is disruptive to traffic. Auto non-DFE mode works in both AN and fixed 8 Gbps speeds.

When used with the **--show** option, the command displays the following information for the specified ports:

Port

The port index number.

8G Non-DFE Configured

Displays ON (Static) if non-DFE is enabled on the port.

Displays OFF (Auto) if the feature is disabled (default).

8G Non-DFE State

Displays Active if non-DFE is activated on the port.

Displays Inactive if non-DFE is not activated.

Use the **portCfgShow** command to display the non-DFE configuration along with other port parameters.

Notes This command is applicable only on an 8 Gbps 16 Gbps-capable port.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|--------------------|--|
| <i>slot/</i> | Specifies the slot number to be configured on bladed systems. |
| <i>port[-port]</i> | Specifies a port or a port range, relative to the slot number on bladed systems. This command does not support multiple port ranges. |
| --enable | Enables non-DFE statically on the specified ports. |
| --disable | Disables non-DFE which allows automatic activation on the specified ports. |
| --show | Displays the non-DFE configuration and state on the specified ports. |
| --help | Displays the command usage. |

Examples To enable non-DFE on a single port:

```
switch:admin> portcfgnondfe --enable 4/28
```

To enable non-DFE on a port range:

```
switch:admin> portcfgnondfe --enable 0-8
```

To enable the non-DFE feature on a range of ports (in this example, non-DFE is already enabled on ports 2 and 3):

```
switch:admin> portcfgnondfe --enable 2-4
Same configuration for port 2
Same configuration for port 3
```

To disable the non-DFE feature on a port range:

```
switch:admin> portcfgnondfe --disable 0-8
```

To display non-DFE configuration on a single port:

```
switch:admin> portcfgnondfe --show 7
Port: 7
8G Non-DFE Configured: OFF (Auto)
8G Non-DFE State: Active
```

```
switch:admin> portcfgnondfe --show 8
Port: 8
8G Non-DFE Configured: OFF (Auto)
```



```
8G Non-DFE State: Inactive
```

```
switch:admin> portcfnondfe --show 9  
Port: 9  
8G Non-DFE Configured: ON (Static)  
8G Non-DFE State: Active
```

SEE ALSO [portCfgShow](#)

portCfgNPIVPort

Enables or disables N_Port ID virtualization (NPIV) functionality on a port and sets the per-port login limit.

Synopsis

```
portcfgnpivport --enable [slot/]port
portcfgnpivport --disable [slot/]port
portcfgnpivport --setloginlimit [slot/]port login_limit
portcfgnpivport --help
portcfgnpivport [slot/]port mode
```

Description Use this command to enable or disable NPIV functionality on a port and to configure the maximum number of logins for this port. Changes made by this command are persistent across switch reboots and power cycles.

N_Port ID Virtualization (NPIV) enables a single Fibre Channel protocol port to appear as multiple, distinct ports, providing separate port identification within the fabric for each operating system image behind the port as if each operating system image had its own unique physical port. NPIV assigns a different virtual port ID to each Fibre Channel protocol device without impacting your existing hardware implementation. The virtual port has the same properties as an N_Port and is therefore capable of registering with all fabric services.

The following conditions must be met for a switch port to respond to NPIV requests from an NPIV device:

NPIV capable

NPIV capability is a switch blade or port attribute that is required for NPIV functionality. Some blades within a switch, or some ports within a switch or blade, may not have NPIV capability. NPIV functionality cannot be enabled on such ports and they do not respond to NPIV requests.

NPIV enabled

NPIV functionality must be enabled on a port for it to respond to NPIV requests. NPIV is enabled by default. It can be selectively disabled or re-enabled on specified switch ports using this command.

NPIV HA

To enable NPIV functionality on dual-CP systems, NPIV-enabled firmware must be running on both the active and the standby CPs. This requirement does not apply to single-CP systems.

Use the **--setloginlimit** option to configure the number of permitted logins per NPIV port. Up to 255 virtual port IDs are allocated per NPIV port, which means each NPIV port can support up to 255 logins. The number of logins you can configure per port is between 0 to 255. The default is 126 per port. The sum of all configured per-port login maximum values cannot exceed the total number of logins set for the switch with the **configure** command.

As of Fabric OS v.6.4.0 you can no longer use the **configure** command to set the NPIV port login limit on a per-switch basis. However, the maximum number of logins per switch is still set and enforced by the **configure** command. Use **configure** to display the current setting.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When using **portCfgNPIVPort --disable** or **--enable** on trunked ports, you must disable all trunk member ports before changing the NPIV capability on a trunk member and then re-enable the trunked ports to ensure that the configuration changes take effect.

Use the **portCfgShow** command to determine whether NPIV is enabled on a port and to display the maximum logins configured for that port. Use the **portCfgDefault** command to reset all port configurations, including the NPIV configuration of a port.

Operands This command has the following operands:

| | |
|------------------------------------|---|
| slot | For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/). |
| port | Specifies the number of the port to be configured, relative to its slot for bladed systems. Use switchShow for a listing of valid ports. |
| --enable | Enables NPIV on the specified port. This command can be executed only in switch mode. |
| --disable | Disables NPIV on the specified port. This command can be executed only in switch mode. |
| --setloginlimit login_limit | Sets the maximum number of allowed logins for the specified port. This command can be executed both in switch mode and Access Gateway mode. Valid values for <i>login_limit</i> are 0 to 255. The default value is 126. |
| --help | Displays the command usage. |
| mode | Enables or disables NPIV on the specified port. Specify 1 to enable or 0 to disable the feature. The mode operand is a legacy command; it will be deprecated in a future Fabric OS release. |

Examples To display the current NPIV port configuration:

```
switch:admin> portcfgshow
Ports of Slot 0   0   1   2   3   4   5   6   7   8   9  10  11  12  13  14  15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed            AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Fill Word(On Active)0 0 0 0  0 0 0 0  0 0 0 0  0 0 0 0
Fill Word(Current)0 0 0 0  0 0 0 0  0 0 0 0  0 0 0 0
Fil 0 0 0 0 0 0 0 0 0 0 0 0 0 0
AL_PA Offset 13  . . . . . . . . . . . . . . . . . . . . . .
Trunk Port       ON ON ON ON  ON ON ON ON  ON ON ON ON  ON ON ON ON
Long Distance    . . . . . . . . . . . . . . . . . . . . . .
VC Link Init     . . . . . . . . . . . . . . . . . . . . . .
Locked L_Port    . . . . . . . . . . . . . . . . . . . . . .
Locked G_Port    . . . . . . . . . . . . . . . . . . . . . .
Disabled E_Port  . . . . . . . . . . . . . . . . . . . . . .
```

```

Locked E_Port      . . . . .
ISL R_RDY Mode    . . . . .
RSCN Suppressed   . . . . .
Persistent Disable.. . ON ON . . . . .
LOS TOV enable    . . . . .
NPIV capability   ON ON ON ON ON ON ON ON ON ON ON ON ON ON
NPIV PP Limit     80 80 80 80 80 80 80 80 95 95 95 95 95 95 95
(output truncated)

```

To disable NPIV functionality on port 7 and to display the change:

```

switch:admin> portcfgnpiport --disable 7
switch:admin> portcfgshow 0/7
Area Number:          7
Speed Level:          AUTO(HW)
Fill Word(On Active): 0(Idle-Idle)
Fill Word(Current):   0(Idle-Idle)
AL_PA Offset 13:     OFF
Trunk Port            ON
Long Distance         OFF
VC Link Init         OFF
Locked L_Port        OFF
Locked G_Port        OFF
Disabled E_Port      OFF
Locked E_Port        OFF
ISL R_RDY Mode       OFF
RSCN Suppressed      OFF
Persistent Disable    OFF
LOS TOV enable       OFF
NPIV capability       OFF
QOS E_Port           AE
Port Auto Disable:   OFF

Rate Limit            OFF
EX Port              OFF
Mirror Port          OFF
Credit Recovery      ON
F_Port Buffers       12
Fault Delay:         0(R_A_TOV)
NPIV PP Limit:       126
CSCTL mode:          OFF
Frame Shooter Port   OFF
D-Port mode:         OFF

```

To increase the maximum logins on port 1 to 255 and to display the change:

```

switch:admin> portcfgnpiport --setloginlimit 1 255

switch:admin> portcfgshow 1
Area Number:          7
Speed Level:          AUTO(HW)
Fill Word(On Active): 0(Idle-Idle)
Fill Word(Current):   0(Idle-Idle)
AL_PA Offset 13:     OFF
Trunk Port            ON
Long Distance         OFF
VC Link Init         OFF
Locked L_Port        OFF
Locked G_Port        OFF

```

| | |
|--------------------|------------|
| Disabled E_Port | OFF |
| Locked E_Port | OFF |
| ISL R_RDY Mode | OFF |
| RSCN Suppressed | OFF |
| Persistent Disable | OFF |
| LOS TOV enable | OFF |
| NPIV capability | OFF |
| QOS E_Port | AE |
| Port Auto Disable: | OFF |
| | |
| Rate Limit | OFF |
| EX Port | OFF |
| Mirror Port | OFF |
| Credit Recovery | ON |
| F_Port Buffers | 12 |
| Fault Delay: | 0(R_A_TOV) |
| NPIV PP Limit: | 255 |
| CSCTL mode: | OFF |
| Frame Shooter Port | OFF |
| D-Port mode: | OFF |

SEE ALSO [configure](#), [portCfgDefault](#), [portCfgShow](#)

portCfgNPort

Enables or disables N_Port functionality for an Access Gateway port.

Synopsis `portcfgnport [slot/]port1[-port2] [mode]`

Description Use this command to enable or disable N_Port functionality for an Access Gateway port or for a range of ports. The enabled N_Ports automatically come online if they are connected to an enterprise fabric switch that supports NPIV. When used without operands, this command displays the port configuration.

Notes NPIV capability must be enabled on the ports connected to the Access Gateway. By default, NPIV is enabled. Use `portcfgnpivport --enable` to enable NPIV capability on a port, if it was previously disabled. Note that enabling or disabling NPIV capability is not supported in Access Gateway mode.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

- slot** Specifies the slot number, followed by a slash (/) on bladed systems.
- port1[-port2]** Specifies a single port or a range of ports to be configured as N_Ports, for example, 3-9, or 2/10-15.
- mode** Enables (1) or disables (0) N_Port functionality on the specified ports. The default mode is 0 (disabled).

Examples To enable N_Port functionality for a port:

```
switch:admin> portcfgnport 8 1
```

To enable N_Port functionality for a set of ports in a specific range:

```
switch:admin> portcfgnport 2-3 1
```

To display the N_Port configuration for all ports:

```
switch:admin> portcfgnport
Ports          0  1  2  3  4  5  6  7  8  9 10  [...]
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Locked N_Port  .. .. ON ON .. .. .. .. ON .. ..  [...]
```

```
switch:admin> portcfgshow
Ports of Slot 0  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed          AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN
Fill Word(On Active)0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

```

Fill Word(Current )0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
Trunk Port          ON ON ON ON ON ON ON ON ON ON ON ON ON ON
Locked N_Port       .. .. ON ON .. .. .. .. .. .. .. ..
Persistent Disable.. .. .. .. .. .. .. .. .. .. .. .. ..
(output truncated)

```

SEE ALSO [portCfgShow, ag](#)

portCfgOctetSpeedCombo

Sets port speed combination for a port octet.

Synopsis `portcfgoctetspeedcombo [slot]port combo`

`portcfgoctetspeedcombo [slot]port -default`

Description

Use this command to configure the speed for a port octet. Each port octet contains eight ports. You can configure any three different speed combinations. When you configure a given port, the combination applies to all ports in the octet. You can specify the octet by any port within the octet. To change the first octet, for example, you can specify any port from 0 through 7 as a port argument value. The following speed combinations are supported:

- Autonegotiated or fixed port speeds of 32Gbps, 16 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps (1)
- Autonegotiated or fixed port speeds of 10 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps (2)
- Autonegotiated or fixed port speeds of 16 Gbps and 10 Gbps (3)

A port octet can be set to any of the three octet combinations, and the ports in the octet can run on any speed supported by its octet combination. This applies to both autonegotiated and fixed speeds.

Before you change the octet speed combination, make sure that the following conditions are met:

- All fixed speed ports are configured at a speed supported by the new combination.
- All online ports in autonegotiation mode have a negotiated speed supported by the new combination.
- If a port is running at a speed not supported in the new combo, you must disable the port or change the speed to a supported fixed speed before you can set the combo.

If any of the ports does not meet the conditions, the operation fails with an appropriate error message. You can change the port speed or disable the ports and retry the command.

The octet combination default is 1. Use the `portCfgOctetSpeedCombo port -default` command to reset the octet combination to its default value.

Notes

This command is supported only on 16 Gbps-capable and 32 Gbps-capable platforms. On the Brocade 6505, the command is not supported.

The operation of the `portCfgOctetSpeedCombo` can be disruptive.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

slot

On bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

port

Sets the speed combination for the specified port and all other ports in the octet port group.

combo

Specifies the speed combination for the ports in the octet. The following speed combinations are supported:

1

Supports autonegotiated or fixed port speeds of 32 Gbps, 16 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps. 2 Gbps speed is not supported on Brocade G620 switch.

2

Supports autonegotiated or fixed port speeds of 10 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps. 2 Gbps speed is not supported on Brocade G620 switch.

3

Supports autonegotiated or fixed port speeds of 16 Gbps and 10 Gbps. Not supported on Brocade G620 switch.

-default

Sets all octet combination values back to the default value of 1.

Examples

To configure the ports in first octet for combination 3 on a Brocade 6510:

```
switch:admin> portcfgoctetspeedcombo 3 3
```

To reset the combination to default:

```
switch:admin> portcfgoctetspeedcombo 3 -default
```

To switch from combo 3 to combo 2 and there are online ports running at 16G speed.:

```
switch:admin> portcfgoctetspeedcombo 2 2
```

The following ports must be disabled or \
speed configuration needs to be change.

Please retry the operation after taking appropriate action.

Speeds supported for octet combo 2 - [10G,8G,4G,2G,AN]

| Port | Speed |
|------|-------|
| 0 | 16* |
| 1 | N16+ |
| 2 | N16+ |
| 3 | N16+ |
| 4 | N10+ |
| 5 | N16+ |
| 6 | N16+ |
| 7 | N16+ |

* Port speed configuration must be changed

+ Port must be disabled

Setting octet speed combo failed

SEE ALSO [portCfgSpeed](#), [switchCfgSpeed](#)

portCfgPersistence

Sets or removes the persistent disable flag on a port or a range of ports.

Synopsis

```
portcfgpersistence --set [-persistentenable | -pe] [-persistentdisable | -pd] [slot]port1[-port2] [...]
```

```
portcfgpersistence --set [-persistentenable | -pe] [-persistentdisable | -pd] -i [index1[-index2] [...] [-f]]
```

```
portcfgpersistence --set [-persistentenable | -pe] [-persistentdisable | -pd] -x [index1[-index2] [...] [-f]]
```

```
portcfgpersistence --set [-persistentenable | -pe] [-persistentdisable | -pd] -slot [slot1[-slot2] [...]
```

```
portcfgpersistence -h
```

Description Use this command to set or remove the persistent disable flag on a port or a range of ports.

You can identify a single port to be configured by its port number or by its port index number. Port ranges are supported with index numbers or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers.

Specifying multiple ports with the index (-i), (-x), or slot (-s) option is supported only if **portSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

This command only sets or removes the flag for persistent disabling of the port. The switch still runs power-on diagnostics and initializes a persistently disabled port. The **portEnable**, **switchEnable**, and **bladeEnable** commands do not enable a specific port or ports alone, but these commands succeed on a switch with one or more persistently disabled ports.

The persistent switch disable or enable configuration does not alter the persistent disable or enable configurations of the ports within the switch.

Because ports are by default persistently enabled, the persistently disabled state of a port is cleared by the **portCfgDefault** command.

Notes This command is blocked if the persistent disable flag is set when the port is currently enabled.

This command is blocked if the switch is operating in the FICON Management Server mode (fmsmode); instead, use **portDisable** with Active=Saved mode enabled.

This command is not supported on FCoE ports. To disable an FCoE port, use **fcoe --disable**.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

On bladed systems only, specifies the slot number of the ports to set the flag, followed by a slash (/).

| | |
|---|--|
| <i>port1</i> [- <i>port2</i>] | Sets the flag for a single port or a range of ports identified by port numbers. The port range cannot span slots, but you can specify multiple port range pairs separated by a space, for example 3/1-4 4/7-9. |
| -f | Ignores nonexistent ports. |
| -i <i>index1</i> [- <i>index2</i>] | Sets the flag for a port or a range of ports identified by port index numbers. You can specify multiple index ranges separated by a space, for example, -i 33-38 40-60. |
| -x <i>index1</i> [- <i>index2</i>] | Sets the flag for a port or a range of ports identified by index number in hexadecimal format. You can specify multiple port ranges separated by a space, for example, -x 21-26 28-3c. |
| -slot [<i>slot1</i> [- <i>slot2</i>] | Sets the flag on all ports on a slot or on a range of slots, for example, -s 3-5. You can specify multiple slot ranges separated by a space, for example, -s 3-5 8-10. |
| -h | Displays the command usage. |

Examples

To set the persistentenable flag on a port:

```
switch:admin> portcfgpersistence --set -pe 2
```

To set the persistentdisable flag on a port:

```
switch:admin> portcfgpersistence --set -pd 0
```

SEE ALSO

[ficonCupSet](#), [ficonCupShow](#), [portCfgDefault](#), [portDisable](#), [portEnable](#), [portCfgPersistentDisable](#), [portCfgPersistentEnable](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portCfgPersistentDisable

Persistently disables a port or a range of ports.

Synopsis

portcfgpersistentdisable

portcfgpersistentdisable [*slot*]*port1*[-*port2*] [...]

portcfgpersistentdisable -i [*index1*[-*index2*] [...] [-f]]

portcfgpersistentdisable -x [*hex1*[-*hex2*] [...] [-f]]

portcfgpersistentdisable -slot [*slot1*[-*slot2*] [...]

portcfgpersistentdisable [*slot*]*port* -r *disable_reason_string*

portcfgpersistentdisable -h

Description

Use this command to persistently disable a port or a range of ports. Persistently disabled ports remain disabled across power cycles, switch reboots, and switch enables. By default, a port is enabled persistently, unless the port is capable of routing. The change in configuration is effective immediately.

You can identify a single port to be configured by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers (decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers. When used without operands, this command displays all persistently disabled ports on the switch.

Specifying multiple ports with the index number (-i or -x) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

The persistent disable configuration overrides existing port configurations, but it does not change these configurations. Use the **portCfgPersistentEnable** command to enable a port persistently and to restore all previously set port configurations for that port. The switch still runs power-on diagnostics and initializes a persistently disabled port. The **portEnable**, **switchEnable**, and **bladeEnable** commands do not enable a specific port or ports alone, but these commands succeed on a switch with one or more persistently disabled ports. The **portEnable** command fails when issued on persistently disabled ports.

The persistent switch disable or enable configuration does not alter the persistent disable or enable configurations of the ports within the switch.

Because ports are by default persistently enabled, the persistently disabled state of a port is cleared by the **portCfgDefault** command.

Notes

This command is blocked if the switch is operating in the FICON Management Server mode (fmsmode); instead, use **portDisable** with Active=Saved mode enabled.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports. To disable an FCoE port, use **fcoe --disable**.

Operands This command has the following operands:

- slot**
On bladed systems only, specifies the slot number of the ports to be disabled persistently, followed by a slash (/).
- port1[-port2]**
Persistently disables a single port or a range of ports identified by port numbers. The port range cannot span slots, but you can specify multiple port ranges pairs separated by a space, for example 3/1-4 4/7-9.
- i index1[-index2]**
Persistently disables a port or a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, -i 33-38 40-60.
- x [hex1 [-hex2]]**
Persistently disables a port or a range of ports identified by port index numbers in hexadecimal format. You may specify multiple port ranges separated by a space, for example, -x 21-26 28-3c.
- f**
Ignores nonexistent ports. This operand is valid only with the **-i** and **-x** options.
- slot [slot1[-slot2]**
Persistently disables all ports on a slot or on a range of slots, for example, **-s 3-5**. You may specify multiple slot ranges separated by a space, for example, **-s 3-5 8-10**.
- r disable_reason_string**
Specifies the reason for disabling the port persistently. The string can be up to 16 characters long and must be enclosed in double quotation marks.
- h**
Displays the command usage.

Examples To disable a single port persistently:

```
switch:admin> portcfgpersistentdisable 2/4
```

To disable a range of ports persistently:

```
switch:admin> portcfgpersistentdisable 2/4-8
```

To disable multiple port ranges persistently:

```
switch:admin> portcfgpersistentdisable 2/12 \
-15 3/10-12 4/3-4
```

To display the **portSwap** status:

```
switch:admin> portcfgpersistentdisable -i
portcfgpersistentdisable: portSwap feature enabled.
'-i' option not supported.
```

To disable a port persistently by specifying its index number:

```
switch:admin> portcfgpersistentdisable -i 176
```

portCfgPersistentDisable

To disable a range of ports persistently by specifying the corresponding port index range:

```
switch:admin> portcfgpersistentdisable -i 17-18
```

To disable multiple ports persistently by specifying multiple port index ranges:

```
switch:admin> portcfgpersistentdisable -i 17-18 30-39
```

To disable all ports on slots 3-5 persistently:

```
switch:admin> portcfgpersistentdisable -s 3-5
```

To disable all ports on slots 3-5 and 8-10 persistently:

```
switch:admin> portcfgpersistentdisable -s 3-5 8-10
```

To display the persistently disabled ports on the switch:

```
switch:admin> portcfgpersistentdisable
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Disabled - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
Slot 0 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Disabled - YES YES - - - - - - - - - - - - - - - - YES YES

Slot 0 32 33 34 35 36 37 38 39
-----+-----+-----+-----+-----+-----+-----+-----+
Disabled YES YES YES YES YES YES YES YES
```

To disable a port persistently with a disable reason:

```
switch:admin> portcfgpersistentdisable 3/3 -r "bad SFP"
```

SEE ALSO

[ficonCupSet](#), [ficonCupShow](#), [portCfgDefault](#), [portDisable](#), [portEnable](#), [portCfgPersistentEnable](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portCfgPersistentEnable

Persistently enables a port or a range of ports.

Synopsis `portcfgpersistentenable`

`portcfgpersistentenable [slot]port1[-port2] [...]`

`portcfgpersistentenable -i [index1[-index2] [...] [-f]]`

`portcfgpersistentenable -x [hex1[-hex2] [...] [-f]]`

`portcfgpersistentenable -s[slot] [slot1[-slot2]`

`portcfgpersistentenable -h`

Description

Use this command to persistently enable a port or a range of ports. If the port is connected to another switch when this command is issued, the fabric may reconfigure. After the port is persistently enabled, devices connected to the port can again communicate with the fabric.

You can identify a single port to be configured by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers (decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers. When used without operands, this command displays all persistently enabled ports on the switch.

Specifying multiple ports with the index number (-i or -x) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

For ports that come online after being persistently enabled, the following indications may be sent to indicate a state transition: RSCN, SNMP trap, Web pop-up window.

Persistently enabled ports remain enabled across power cycles, switch reboots, and switch enables. By default, a port is enabled persistently, unless the port is capable of routing. The change in configuration is effective immediately.

This command re-enables all previously set port configurations of a specified port. You can temporarily disable a persistently enabled port with the **portDisable** or **switchDisable** commands. The persistent switch disable or enable configuration does not alter the persistent disable or enable configurations of the ports within the switch. The configuration commands **configDefault** and **portCfgDefault** do not modify the persistent enable attribute of a port.

Notes

This command is blocked if the switch is operating in the FICON Management Server mode (fmsmode). Instead use **portEnable** with Active=Saved Mode enabled.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports. To enable an FCoE port, use **fcoe --enable**.

Operands This command has the following operands:

| | |
|---------------------------------------|--|
| <i>slot</i> | On bladed systems only, specifies the slot number of the ports to be enabled persistently, followed by a slash (/). |
| <i>port1[port2]</i> | Persistently enables a single port or a range of ports identified by port numbers. The port range cannot span slots, but you can specify multiple port ranges pairs separated by a space, for example 3/1-4 4/7-9. |
| -i <i>index1[-index2]</i> | Persistently enables a port or a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, 33-38 40-60. |
| -x [<i>hex1 [-hex2]</i>] | Persistently enables a port or a range of ports identified by port index numbers in hexadecimal format. You may specify multiple port ranges separated by a space, for example, -x 21-26 28-3c. |
| -f | Ignores nonexistent ports. This operand is valid only with the -i and -x options. |
| -slot [<i>slot1[-slot2]</i>] | Persistently enables all ports on a slot or on a range of slots, for example, -s 3-5. Multiple slot ranges are not supported with this command. |
| -h | Displays the command usage. |

Examples To enable a single port persistently:

```
switch:admin> portcfgpersistentenable 2/4
```

To enable a range of ports persistently:

```
switch:admin> portcfgpersistentenable 2/4-8
```

To enable multiple port ranges persistently:

```
switch:admin> portcfgpersistentenable 2/24-26 3/10-12 4/3-4
```

To enable a port persistently by specifying its index number:

```
switch:admin> portcfgpersistentenable -i 176
```

To enable a range of ports persistently by specifying the corresponding port index range:

```
switch:admin> portcfgpersistentenable -i 17-18
```

To enable multiple ports persistently by specifying multiple port index ranges:

```
switch:admin> portcfgpersistentenable -i 17-18 30-39
```


To enable all ports on slots 3-5 persistently:

```
switch:admin> portcfgpersistentenable -s 3-5
```

To display the persistently enabled ports on the switch:

```
switch:admin> portcfgpersistentenable
```

```
Slot 9  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---
Enabled YES YES YES YES YES YES YES YES YES YES YES YES YES YES YES YES
(output truncated)
```

SEE ALSO [ficonCupSet](#), [ficonCupShow](#), [portDisable](#), [portEnable](#), [portCfgPersistentDisable](#), [portCfgShow](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portCfgQoS

Enables or disables QoS, sets the default configuration, and sets and resets the ingress rate limit.

Synopsis

```
portcfgqos --disable | --enable [slot/]port
portcfgqos --setratelimit [slot/]port ratelimit
portcfgqos --resetratelimit [slot/]port
portcfgqos --enable | --disable [slot/]port[-port] csctl_mode
portcfgqos --default [slot/]port
portcfgqos --help
```

Description Use this command to configure traffic prioritization on a port. Two alternate modes of traffic prioritization are supported by this command: Adaptive Networking/Quality of Service (AN/QoS) or Class-Specific Control (CS_CTL):

- The Adaptive Networking with QoS feature allows latency-sensitive applications to share storage resources alongside throughput-intensive applications. You can enable or disable Adaptive Networking/Quality of Service (AN/QoS) on a port, set or reset the ingress rate limit for the specified port, and set the default behavior. Ingress Rate Limiting delays the return of BB credits to the external device. By limiting the throughput on the ingress side of a port, existing congestion can be removed or avoided.
- In Fabric OS v7.0.0 and later, an alternate method of traffic flow prioritization based on the CS_CTL bits of a Fibre Channel frame is provided through this command. This feature uses the value of the CS_CTL bit of the frame to determine the virtual channel (VC), so each frame can be prioritized based on the value of the CS_CTL bit.

CS_CTL flow prioritization is independent of traffic prioritization based on QoS zones; and both methods are mutually exclusive. If CS_CTL VC mode is enabled on an F/FL_Port, QoS-based traffic flow prioritization cannot be used between any two devices connected to these F/FL_Ports in that fabric and vice versa.

If both QoS-based and CS_CTL-based traffic prioritization are enabled on the same F/FL_Port, the CS_CTL-based method takes priority over the QoS zones. When QoS is enabled on an F/FL_Port and you enable CS_CTL VC mode on the same port, the command displays a message stating that QoS zones will lose priority to CS_CTL-based traffic prioritization. When you disable CS_CTL mode on a given F/FL port, the QoS zones, if already enabled, become the effective frame classification method for all devices connected to that F/FL_Port.

On 16 Gbps-capable Inter Chassis Link (ICL) ports, QoS is not configurable and it is always enabled internally.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Enabling and disabling QoS is potentially disruptive to the I/O on the affected port.

This command is not supported on all platforms. Refer to the *Fabric OS Administrator's Guide* for more information.

The configuration changes effected by this command are persistent across system reboots.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to configure, followed by a slash (/).

port[-port]

Specifies the port or port range to be configured, relative to the slot for bladed systems. Use **switchShow** for a listing of valid ports.

--disable

Disables the current configuration on the specified ports. When issued with the **csctl_mode** operand, this command disables traffic prioritization based on CS_CTL. If QoS was enabled before you enabled **csctl_mode**, this command restores QoS-based traffic prioritization.

--enable

Enables QoS or CS_CTL mode. When issued with the **csctl_mode** operand, this command enables traffic prioritization based on CS_CTL. When you enable **csctl_mode**, the command displays a message stating that QoS zones will lose priority to CS_CTL-based traffic prioritization.

--resetratelimit

Turns off the ingress rate limiting feature on the specified ports.

--setratelimit *ratelimit*

Sets an ingress rate limit to reduce traffic from the specified port. This configuration is applicable only to F/FL_Ports. For E/EX_Ports, this configuration would not be effective. The ingress rate limit is enforced only when a given port can run at a speed higher than the speed specified in the configuration. For example if the rate limit is set at 4 Gbps and the port comes online only at 2 Gbps, no enforcement is needed. Specify an ingress rate in Mbps. Supported values for *ratelimit* are: 200, 400, 600, 800, 1000, 1500, 2000, 2500, 3000, 3500, 4000, 5000, 6000, 7000, and 8000, 9000, 10000, 11000, 12000, 13000, 14000, 15000, 16000. This operand is not supported on Brocade G620 switch, Brocade X6-4 Director, and Brocade X6-8 Director.

--default

Applies the default configuration to the specified port. This command attempts to re-enable QoS; success depends on availability of buffers. If CS_CTL was enabled on the port, this command disables the configuration.

--help

Displays the command usage.

Examples

To enable QoS on a range of ports:

```
switch:admin> portcfgqos --enable 1/15-18
```

To disable QoS on a port:

```
switch:admin> portcfgqos --disable 1/15
```

To set the default QoS configuration on a port:

```
switch:admin> portcfgqos --default 12/41
```

To set the ingress rate limit on a port to 2 Gbps:

```
switch:admin> portcfgqos --setratelimit 12/41 2000
```

To reset the ingress rate limit on a port:

```
switch:admin> portcfgqos --resetratelimit 12/41
```

To enable CS_CTL VC mode on a port when QoS is enabled:

```
switch:admin> portcfgqos --enable 1/10-16 csctl_mode
Enabling CSCTL mode flows causes QoS zone flows to \
lose priority on such ports
Do you want to proceed?(y/n): y
```

```
switch:admin> portcfgshow -slot 1
Index:          1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed           AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN
Fill Word(On Active)0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
Fill Word(Current) 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
AL_PA Offset 13  . . . . . . . . . . . . . . . . . . . . . . . .
Trunk Port      ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
Long Distance   . . . . . . . . . . . . . . . . . . . . . . . .
VC Link Init    . . . . . . . . . . . . . . . . . . . . . . . .
Locked L_Port   . . . . . . . . . . . . . . . . . . . . . . . .
Locked G_Port   . . . . . . . . . . . . . . . . . . . . . . . .
Disabled E_Port . . . . . . . . . . . . . . . . . . . . . . . .
Locked E_Port   . . . . . . . . . . . . . . . . . . . . . . . .
ISL R_RDY Mode  . . . . . . . . . . . . . . . . . . . . . . . .
RSCN Suppressed . . . . . . . . . . . . . . . . . . . . . . . .
Persistent Disable . . . . . . . . . . . . . . . . . . . . . . . .
LOS TOV enable  . . . . . . . . . . . . . . . . . . . . . . . .
NPIV capability ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
NPIV PP Limit   126 126 126 126 126 126 126 126 126 126 126 126 126 126 126 126
QOS Port        AE AE AE AE AE AE AE AE AE AE AE AE AE AE AE AE AE
EX Port         ON ON . . . . . . . . . . . . . . . . . . . . . .
Mirror Port     . . . . . . . . . . . . . . . . . . . . . . . .
Rate Limit      . . . . . . . . . . . . . . . . . . . . . . . .
Credit Recovery ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
Fport Buffers   . . . . . . . . . . . . . . . . . . . . . . . .
Port Auto Disable . . . . . . . . . . . . . . . . . . . . . . . .
CSCTL mode      . . . . . . . . . . . . . . . . . . ON ON ON ON ON ON ON ON
D-Port mode     . . . . . . . . . . . . . . . . . . . . . . . .
Fault Delay     0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
(output truncated)
```

To disable CS_CTL VC mode on a port range:

```
switch:admin> portcfgqos --disable 1/10-16 csctl_mode
```

SEE ALSO [portCfg](#), [portCmd](#), [portShow](#), [switchShow](#), [configure](#)

portCfgShow

Displays port configuration settings.

Synopsis

```
portcfgshow
portcfgshow [slot]port
portcfgshow -i [index1[-index2] [...] [-f]]
portcfgshow -slotslot1[-slot2] [...]
portcfgshow [option] [slot]ge_port
portcfgshow [fciptunnel | fcipcircuit]
    [all | [slot]veport] [optional_argument]
portcfgshow ipsec-policy
portcfgshow lag
portcfgshow tcl
portcfgshow sla
```

Description

Use this command to display the current configuration of a port. The behavior of this command is platform-specific; output varies depending on port type and platform, and not all options are supported on all platforms.

Non-GbE port displays

When used without operands, this command displays port configuration settings for all ports on a switch, except Gigabit Ethernet (GbE) ports.

Displays supported only on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades

Use this command with optional arguments to display the following parameters configured for a GbE port on the Brocade 7840 switch and on the Brocade FX8-24 and Brocade SX6 blades.

- IP interfaces on the GbE port
- Static routes on the IP interfaces
- Address resolution protocol (ARP) entries
- VLAN tag configuration
- Fibre Channel over IP (FCIP) tunnel configuration settings
- IPSec configuration
- Inband Management IP addresses and routes.
- Display IPsec policy. Currently supported on the Brocade 7840 and Brocade SX6 blade only.

Non-GbE port displays

The following information is displayed when the command is issued for all ports, for a single port, or for one or more port ranges specified by their port index numbers or slot numbers:

Area Number

Displays the port area number. This field is displayed only when **portCfgShow** is executed for a single port.

Octet Speed Combo

Displays the speed configuration for a port octet. This value is set by the **portCfgOctetSpeedCombo** command. Port octet speed configuration is supported only on 16 Gbps-capable platforms. On unsupported platforms the Octet Speed Combo field is suppressed. Valid Octet Combo values include the following:

1

Autonegotiated or fixed port speeds of 16 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps.

2

Autonegotiated or fixed port speeds of 10 Gbps, 8 Gbps, 4 Gbps, and 2 Gbps.

3

Autonegotiated or fixed port speeds of 16 Gbps and 10 Gbps.

Speed

Displays Auto for auto speed negotiation mode, or a specific speed of 1 Gbps, 2 Gbps, 4 Gbps, 8 Gbps, or 16 Gbps, and soft, 10, or AX on 8 Gbps. This value is set by the **portCfgSpeed** command.

Fill Word (On Active)

Displays the fill word setting that will take effect after the next time the port goes offline and comes back. Possible values are 0(Idle-Idle), 1(arbff-arbff), 2(Idle-arbff), or 3(A-A then SW I-A). This parameter is set by the **portcfgFillword** command.

Fill Word (Current)

Displays the currently effective fill word setting. Possible values are 0(Idle-Idle), 1(arbff-arbff), 2(Idle-arbff), or 3(A-A then SW I-A). This parameter is set by the **portcfgFillword** command.

AL_PA Offset 13

Displays (..) or OFF when the arbitrated loop physical address (AL_PA) on the port is configured to use a 0x0 AL_PA address (default). Displays ON when the address configuration is 0x13 AL_PA. This value is set by the **portCfgAlpa** command.

Trunk Port

Displays ON when port is set for trunking. Displays (..) or OFF when trunking is disabled on the port. This value is set by the **portCfgTrunkPort** command.

Long Distance

Displays (..) or OFF when long distance mode is off; otherwise, displays long distance levels. This value is set by the **portCfgLongDistance** command. Values are as follows:

LE

The link is up to 10 km.

LD

The distance is determined dynamically.

LS

The distance is determined statically by user input.

VC Link Init

Displays (..) or OFF when the long distance link initialization option is turned off. Displays ON when it is turned on for long distance mode. This value is set by the **portCfgLongDistance** command.

Desired Distance

Displays the desired distance of the port. This field is displayed only when portCfgShow is executed for a single port and would only display for LS and LD long distance mode.

Reserved Buffers

Displays the reserved buffers for the port. This field is displayed only when portCfgShow is executed for a single port and would only display for LS and LD long distance mode.

Locked L_Port

Displays ON when the port is locked to L_Port only. Displays (..) or OFF when L_Port lock mode is disabled and the port behaves as a U_Port). This value is set by the **portCfgLport** command.

Locked G_Port

Displays ON when the port is locked to G_Port only. Displays (..) or OFF when G_Port lock mode is disabled and the port behaves as a U_Port. This value is set by the **portCfgGport** command.

Disabled E_Port

Displays ON when the port is not allowed to be an E_Port. Displays (..) or OFF when the port is allowed to function as an E_Port. This value is set by the **portCfgEport** command.

ISL R_RDY Mode

Displays ON when ISL R_RDY mode is enabled on the port. Displays (..) or OFF when ISL R_RDY mode is disabled. This value is set by the **portCfgISLMode** command.

| | |
|---------------------------|--|
| RSCN Suppressed | Displays ON when RSCN suppression is enabled on the port. Displays (..) or OFF when RSCN suppression is disabled. This value is set by the portCfg rscnsupr command. |
| Persistent Disable | Displays ON when the port is persistently disabled; otherwise displays (..) or OFF. This value is set by the portCfgPersistentDisable command. |
| LOS TOV mode | Displays 1 (ON:FixSpd) or 2 (ON:AN&FixSpd) when LOS TOV is enabled on the port; otherwise displays (..) or 0 (OFF). This value is set by the portCfgLossTov command. |
| NPIV capability | Displays ON when N_Port ID Virtualization (NPIV) is enabled on the port (default). Displays (..) or OFF when NPIV capability is disabled. This value is set by the portCfgNPIVPort command. |
| NPIV PP Limit | Displays the maximum number of allowed logins for the port. Displays the default of 126 or the configured maximum. This parameter is set with the portCfgNPIVPort --setloginlimit command. |
| NPIV FLOGI Logout | Displays ON when the Base Device Logout feature is enabled and the base device can logout without disrupting the NPIV devices on the same port. Displays (..) when the base device log out causes the NPIV devices on the same port to log out. |
| QoS E_Port | <p>Displays ON when Quality of Service (QoS) is enabled on the E_Port (or EX_Port) when QoS is enabled in an FCR deployment scenario. Displays (..) or OFF when QoS is disabled. By default, QoS is enabled if sufficient buffers are available. Displays AE when QoS is configured as Auto Enabled. In the AE state, QoS is enabled based on the availability of buffers.</p> <p>Use islShow to determine the current status of QoS (on or off) in the AE state. This value is set by the portCfgQos command.</p> |
| EX_Port | Displays ON when the port is configured as an EX_Port. Otherwise displays (..) or OFF. This value is set by the portCfgExPort command. |
| Mirror Port | Displays ON when Mirror Port is enabled on the port. Displays (..) or OFF when Mirror Port is disabled. This value is set by the portCfg mirrorport command. |
| SIM Port | Displays ON when SIM Port is enabled on the port. Displays (..) or OFF when SIM Port is disabled. This value is set by the flow --control flow_name -simport command. |
| FC Fastwrite | Displays ON when FC FastWrite is enabled on the port or (..) or OFF when disabled. FastWrite is disabled by default. This value is set by the portCfg fastwrite FC Fastwrite is no longer supported as of Fabric OS v7.0.0. |

| | |
|--------------------------|--|
| Rate Limit | Displays ON when ingress rate limit is set on the port or (..) or OFF when the ingress rate limiting feature is disabled. This value is set by the portCfgQos --setratelimit command. The default value is OFF. |
| Credit Recovery | Displays ON when Credit Recovery is enabled on the port. Displays (..) or OFF when the feature is disabled. This value is set by the portCfgCreditRecovery command. The credit recovery feature is enabled by default. |
| Port Auto Disable | Displays On when the Auto Disable feature is enabled on a port. Displays (..) or OFF when the feature is disabled. This feature causes ports to become disabled when they encounter an event that would cause them to reinitialize. This feature is enabled by the portCfgAutoDisable command. The feature is disabled by default. In the single port view, the configured trigger conditions are displayed when the feature is disabled. Refer to the example section for an illustration. |
| F_Port Buffers | Displays the number of configured F_Port buffers. Displays (..) or OFF if no buffers are configured. The buffer value is set by the portCfgfPortbuffers command. |
| E_Port Credits | Displays the number of configured E_Port credits. Displays (..) or OFF if no credits are configured. The credit value is set by the portCfgEportCredits command. |
| CSCTL mode | Displays ON if CSCTL mode is enabled on the port. Displays (..) or OFF if the feature is disabled. This parameter is set with the portCfgQos command. |
| TDZ mode | Displays ON when Target Driven Zoning (TDZ) is enabled on a port and the port is online or offline. Displays (..) or OFF when the configuration is disabled. TDZ is enabled by the portCfgTdz command; it is disabled by default. |
| D_Port mode | Displays ON when the port is configured as a D_Port; otherwise displays (..) or OFF. Refer to the portCfgDPort command for more information. This parameter is displayed only on on 16 Gbps-capable blades that support D_Port capability. |
| D_Port over DWDM | Displays ON when dwdm is enabled on that port using the portcfgdport command. This option is used to set if the link connecting two switches is a DWDM (Dense Wavelength Division Multiplexing) link. However, it can be set on normal links also. |
| Fault Delay | Displays the Fault delay value. Displays 0 if the value is R_A_TOV. This is the default. Displays 1 if the value is 1.2 seconds. This value is set by the portCfgFaultDelay command. valid for FC ports only. |
| Compression | Displays ON when compression is enabled on a port. Displays (..) or OFF when the configuration is disabled. Compression is enabled by the portCfgCompress command; it is disabled by default. |

Encryption

Displays ON when encryption is enabled on a port. Displays (..) or OFF when the configuration is disabled. Encryption is enabled by the **portCfgEncrypt** command; it is disabled by default.

10G/16G FEC

Displays ON when Forward Error Correction (FEC) is enabled on a port and the port is online. Displays (..) or OFF when the configuration is disabled. FEC is enabled by the **portCfgFec** command; it is enabled by default.

16G FEC via TTS

Displays ON when control of the Forward Error Correction (FEC) state is permitted via TTS by an externally attached host or device. Displays (..) or OFF when the external control of FEC is disabled. Refer to the **portCfgFec** command for additional information.

When **portCfgShow** is issued for a single FCoE port, only the following fields are displayed: **Area Number**, **Speed Level**, **AL_PA Offset**, **RSCN Suppressed**, and **NPIV PP Limit**. If the command is issued without operands on a DCX with one or more FCOE10-24 blades, the port parameters that are not valid on the FCoE ports are marked with a dash (-). Refer to the example section for an illustration.

You can identify a single port to be displayed by its port number or by its port index number. Port ranges are supported with index numbers or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers. When used without operands, this command displays all persistently disabled ports on the switch.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The Fabric OS port configuration commands are not supported on FCoE ports.

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

port

Specifies the number of the port to be displayed, relative to its slot for bladed systems. Use **switchShow** for a listing of valid port numbers.

-i index1[-index2]

Specifies a port or a range of ports identified by port index numbers. You can specify multiple index ranges separated by a space, for example, -i 33-38 40-60.

-f

Ignores nonexistent ports. This operand is valid only with the -i option.

-slot slot1[-slot2]

Specifies a slot or a range of slots. You can specify multiple slot ranges separated by a space, for example, -s 1-3 5-7.

ge_port

Specifies the number of the GbE Port to be displayed. The GbE ports are numbered ge0 - ge9 on the Brocade FX8-24 blade. The 10GbE ports on the Brocade FX8-24 blade are numbered xge0 and xge1. Use the **switchShow** command for a list of valid ports.

Use **portCfgShow** with a GbE port or with one of the optional arguments to display specific FCIP-related parameters.

Displays supported on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades

When issued on the Brocade 7840 switch, and Brocade FX8-24 and Brocade SX6 blades, tunnels and parameters not applicable to these platforms are not displayed. Use the **portShow** command to display FCIP tunnel and circuit parameters on the Brocade FX8-24 blade.

ipif

Displays the IP interface configurations. IPv6 addresses are supported.

arp

Displays the address resolution protocol (ARP) table. This option is supported on Brocade FX8-24 and Brocade SX6 extension blades.

iproute

Displays the IP route on the specified GbE port. IPv6 addresses are supported.

vlantag

Displays the VLAN tagging configuration. For each entry, the output displays the IP interface address, the destination IP address, the VLAN ID, and the L2 CoS priority. This display includes permanent entries only. Permanent entries are configured at the IP interface level with the **portCfg vlantag** command. To view VLAN tagged tunnels and circuits along with permanent entries, use the **portShow vlantag** command. This option is supported on Brocade FX8-24 platforms

mgmtif ge_port

Displays the inband management interfaces configured for a specified GbE Port. The display includes the interface status (enabled or disabled), the interface IPv4 Address, the netmask, effective MTU, and annotated port flags. This option is supported on Brocade FX8-24 and Brocade SX6 extension blades.

mgmtroute ge_port

Displays the management routes configured for a specified GbE Port. The display includes the destination IPv4 Address, the netmask, the gateway address, and annotated port flags. This option is supported on Brocade FX8-24 and Brocade SX6 extension blades

fcipcircuit

Displays FCIP circuits and related parameters.

fciptunnel

Displays FCIP tunnels and related parameters.

The following options are supported with **fcipcircuit** and **fciptunnel**:

all

Displays information for all FCIP tunnels.

ve_port

Displays information for the specified FCIP tunnel. To display the tunnel, specify the VE_Port number associated with the tunnel configured on one of the GbE ports. VE_Ports are numbered 16-23.

optional_argument

The following optional argument is supported with **fciptunnel**:

-ipsec

Displays whether IPsec is enabled or disabled. If enabled, the key is displayed. If IPsec is enabled and configured in legacy mode, the mode is displayed as "legacy". The mode information is displayed whether or not you issue the **-ipsec** option.

ipsec-policy

Displays the IPsec policy. This option is supported on the Brocade 7840 switch only.

lag

Displays the static link aggregation group (LAG) information.

tcl

Displays the Traffic Control List (TCL) information.

sla

Displays the Service Level Agreement (SLA) information.

Examples

To display the port configuration settings for a single port on a Brocade DCX 8510-8:

```
switch:admin> portcfgshow 10/38
Area Number:                222
Octet Speed Combo:          1(16G|8G|4G|2G)
Speed Level:                 AUTO(HW)
AL_PA Offset 13:            OFF
Trunk Port                   ON
Long Distance                LS
VC Link Init                 OFF
Desired Distance             10 Km
Reserved Buffers             86
Locked L_Port                OFF
Locked G_Port                OFF
Disabled E_Port              OFF
Locked E_Port                OFF
ISL R_RDY Mode               OFF
RSCN Suppressed              OFF
Persistent Disable           OFF
LOS TOV mode                 2(ON:AN&FixSpd)
NPIV capability              ON
NPIV FLOGI Logout            ON
QOS E_Port                   AE
Port Auto Disable:           ON   LOSG   OLS   NOS   LIP
Rate Limit                   OFF
EX Port                      OFF
Mirror Port                  OFF
SIM Port                     ON
Credit Recovery              ON
F_Port Buffers               OFF
E_Port Credits               OFF
Fault Delay:                 0(R_A_TOV)
NPIV PP Limit:               126
```

```

CSCTL mode:          OFF
TDZ mode:            OFF
D-Port mode:        ON
D-Port over DWDM:   ON
Compression:        OFF
Encryption:          OFF
10G/16G FEC:        ON
16G FEC via TTS:    ON

```

To display the port configuration settings on a Brocade 7840 for a range of ports specified by their index numbers:

```

switch:admin> portcfgshow -i 24-31
Index:          24 25 26 27 28 29 30 31
-----+-----+-----+-----+-----+-----+-----+-----
Speed          AN  AN  AN  AN  AN  AN  AN  AN
Fill Word(On Active) 0  0  0  0  0  0  0  0
Fill Word(Current)  0  0  0  0  0  0  0  0
AL_PA Offset 13  ..  ..  ..  ..  ..  ..  ..  ..
Trunk Port      ..  ..  ..  ..  ..  ..  ..  ..
Long Distance   ..  ..  ..  ..  ..  ..  ..  ..
VC Link Init    ..  ..  ..  ..  ..  ..  ..  ..
Locked L_Port   ..  ..  ..  ..  ..  ..  ..  ..
Locked G_Port   ..  ..  ..  ..  ..  ..  ..  ..
Disabled E_Port ..  ..  ..  ..  ..  ..  ..  ..
Locked E_Port   ..  ..  ..  ..  ..  ..  ..  ..
ISL R_RDY Mode  ..  ..  ..  ..  ..  ..  ..  ..
RSCN Suppressed ..  ..  ..  ..  ..  ..  ..  ..
Persistent Disable .. .. .. .. .. .. .. ..
LOS TOV mode    0  0  0  0  0  0  0  0
NPIV capability ON  ON  ON  ON  ON  ON  ON  ON
NPIV PP Limit   126 126 126 126 126 126 126 126
NPIV FLOGI Logout .. .. .. .. .. .. .. ..
QOS Port        ..  ..  ..  ..  ..  ..  ..  ..
EX Port         ..  ..  ..  ..  ..  ..  ..  ..
Fport Buffers   ..  ..  ..  ..  ..  ..  ..  ..
Eport Credits   ..  ..  ..  ..  ..  ..  ..  ..
Port Auto Disable .. .. .. .. .. .. .. ..

Fault Delay     0  0  0  0  0  0  0  0
SIM Port        ..  ..  ..  ..  ..  ..  ..  ..
TDZ mode        ..  ..  ..  ..  ..  ..  ..  ..

```

where AE:QoSAutoEnable, AN:AutoNegotiate, ..:OFF,
-:NotApplicable, ??:INVALID,

To display the port configuration settings for all ports on slot 1 on a Brocade DCX 8510-4:

```

switch:admin> portcfgshow -slot 1
Index:          0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
Octet Speed Combo 1  1  1  1  1  1  1  1  1  1  1  1  1  1  1  1
Speed            AN  AN  AN  AN  AN  AN  AN  AN  AN  AN  AN  AN  AN  AN  AN  AN
AL_PA Offset 13  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..
Trunk Port        ON  ON  ON  ON  ON  ON  ON  ON  ON  ON  ON  ON  ON  ON  ON  ON
Long Distance     ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..
VC Link Init      ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..
Locked L_Port     ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..
Locked G_Port     ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..
Disabled E_Port   ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..

```



```

Fault Delay      0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
SIM Port         .. .. .. .. .. .. .. .. .. .. .. .. .. ..

```

where AE:QoSAutoEnable, AN:AutoNegotiate, ..:OFF, NA:NotApplicable,
 ??:INVALID,

To display the configuration settings for a single port on a switch with Access Gateway enabled:

```

switch:admin> portcfgshow 8
Area Number:      8
Speed Level:      AUTO(HW)
Fill Word(On Active) 0(Idle-Idle)
Fill Word(Current) 0(Idle-Idle)
Trunk Port        ON
Locked N_Port     OFF
Persistent Disable OFF
NPIV capability   ON
NPIV FLOGI Logout ON
QOS Port          AE
Port Auto Disable: OFF
Rate Limit        OFF
F_Port Buffers   OFF
NPIV PP Limit:    126
Fault Delay:      0(R_A_TOV)
NPIV PP Limit:    126
D-Port mode:      OFF
Compression:      OFF
Encryption:       OFF
10G/16G FEC:      ON
16G FEC via TTS:  ON

```

To display port configuration settings on a Gen6 switch with a Brocade SX6 blade in Slot 4:

```

switch:admin> portcfgshow -slot 12
Index:           64 65 66 67 68 69 70 71 [truncated]
-----+-----+-----+-----+-----+-----+-----+-----+-----+
Octet Speed Combo  1  1  1  1  1  1  1  1
Speed              AN  AN  AN  AN  AN  AN  AN  AN
AL_PA Offset 13    .. .. .. .. .. .. .. ..
Trunk Port         ON  ON  ON  ON  ON  ON  ON  ON
Long Distance      .. .. .. .. .. .. .. ..
VC Link Init       .. .. .. .. .. .. .. ..
Locked L_Port      .. .. .. .. .. .. .. ..
Locked G_Port      .. .. .. .. .. .. .. ..
Disabled E_Port    .. .. .. .. .. .. .. ..
Locked E_Port      .. .. .. .. .. .. .. ..
ISL R_RDY Mode     .. .. .. .. .. .. .. ..
RSCN Suppressed    .. .. .. .. .. .. .. ..
Persistent Disable  .. .. .. .. .. .. .. ..
LOS TOV mode       0  0  0  0  0  0  0  0
NPIV capability    ON  ON  ON  ON  ON  ON  ON  ON
NPIV PP Limit      126 126 126 126 126 126 126 126
NPIV FLOGI Logout  .. .. .. .. .. .. .. ..
QOS Port           AE  AE  AE  AE  AE  AE  AE  AE
EX Port            .. .. .. .. .. .. .. ..
Mirror Port        .. .. .. .. .. .. .. ..
Credit Recovery    ON  ON  ON  ON  ON  ON  ON  ON
Fport Buffers      .. .. .. .. .. .. .. ..

```

portCfgShow

```
Eport Credits      .. .. .. .. .. .. .. ..
Port Auto Disable  .. .. .. .. .. .. .. ..
CSCTL mode         .. .. .. .. .. .. .. ..
D-Port mode       .. .. .. .. .. .. .. ..
D-Port over DWDM  .. .. .. .. .. .. .. ..
Compression        .. .. .. .. .. .. .. ..
10G/16G FEC        ON ON ON ON ON ON ON ON
16G FEC via TTS    .. .. .. .. .. .. .. ..
Fault Delay        0 0 0 0 0 0 0 0
SIM Port           .. .. .. .. .. .. .. ..
TDZ mode           .. .. .. .. .. .. .. ..
[output truncated]
```

To display port configuration settings for a single port on a Gen6 switch with a Brocade SX6 blade:

```
switch:admin> portcfgshow 4/4
Area Number:          68
Octet Speed Combo:    1(32G|16G|8G|4G)
Speed Level:          AUTO(HW)
AL_PA Offset 13:      OFF
Trunk Port            ON
Long Distance         OFF
VC Link Init          OFF
Locked L_Port         OFF
Locked G_Port         OFF
Disabled E_Port       OFF
Locked E_Port         OFF
ISL R_RDY Mode        OFF
RSCN Suppressed       OFF
Persistent Disable    OFF
LOS TOV mode          0(OFF)
NPIV capability        ON
QOS Port              AE
Port Auto Disable:    OFF
EX Port               OFF
Mirror Port           OFF
SIM Port              OFF
Credit Recovery       ON
F_Port Buffers        OFF
E_Port Credits        OFF
Fault Delay:          0(R_A_TOV)
NPIV PP Limit:        126
NPIV FLOGI Logout:    OFF
CSCTL mode:           OFF
TDZ mode:             OFF
D-Port mode:          OFF
D-Port over DWDM:    OFF
Compression:          OFF
10G/16G FEC:         ON
16G FEC via TTS:     OFF
```


To display the QoS configuration for an EX_Port (QoS over FCR deployment):

```
switch:admin> switchshow | grep EX-Port
 16 16 id N4 Online FC EX-Port \
    10:00:00:05:1e:41:4a:45 "Tom_100"
(fabric id = 25 )(Trunk master)

switch:admin> portcfgshow 16
Area Number: 11
Octet Speed Combo: 1(16G|8G|4G|2G)
Speed Level: AUTO(HW)
Fill Word(On Active) 0(Idle-Idle)
Fill Word(Current) 0(Idle-Idle)
AL_PA Offset 13: OFF
Trunk Port ON
Long Distance LS
VC Link Init OFF
Desired Distance 10 Km
Reserved Buffers 86
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF
LOS TOV mode 0(OFF)
NPIV capability ON
NPIV FLOGI Logout ON
QOS E_Port AE
Port Auto Disable: OFF
Rate Limit OFF
EX Port ON
Mirror Port OFF
Credit Recovery ON
F_Port Buffers OFF
NPIV PP Limit: 126
CSCTL mode: OFF
TDZ mode: OFF
Fault Delay 1(1.2sec)
D-Port mode: ON
D-Port over DWDM: ON
10G/16G FEC: ON
16G FEC via TTS: ON
```

To display an FCIP Tunnel on a Brocade FX8-24 with and without IPsec information (In the following examples, IPsec is in legacy mode):

```
switch:admin> portcfgshow fciptunnel 1/12 -ipsec
-----
Tunnel ID: 1/12
Tunnel Description:
Compression: Off
Fastwrite: Off
Tape Acceleration: Off
TPerf Option: Off
IPSec: Enabled (legacy)
IPSec Key: '12345678901234567890123456789012'
QoS Percentages: High 50%, Med 30%, Low 20%
Remote WWN: Not Configured
```

portCfgShow

```
Local WWN: 10:00:00:05:1e:52:fe:00  
Flags: 0x00000000  
FICON: Off
```

```
switch:admin> portcfgshow fcipunnel 1/12
```

```
-----  
Tunnel ID: 1/12  
Tunnel Description:  
Compression: Off  
Fastwrite: Off  
Tape Acceleration: Off  
TPerf Option: Off  
IPSec: Enabled (legacy)  
QoS Percentages: High 50%, Med 30%, Low 20%  
Remote WWN: Not Configured  
Local WWN: 10:00:00:05:1e:52:fe:00  
Flags: 0x00000000  
FICON: Off
```

SEE ALSO

[portCfg](#), [portCfgAutoDisable](#), [portCfgCreditRecovery](#), [portCfgEport](#), [portCfgFec](#), [portCfgGport](#), [portCfgLongDistance](#), [portCfgLport](#), [portCfgNPiVPort](#), [portCfgOctetSpeedCombo](#), [portCfgSpeed](#), [portCfgTrunkPort](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portCfgSpeed

Configures the speed for a port or a range of ports.

Synopsis `portcfgspeed [slot/]port speed`

`portcfgspeed -i index1[-index2] [-f] speed [-m] max_auto_speed`

`portcfgspeed -x hex1[-hex2] [...] speed`

`portcfgspeed -slot slot1[-slot2] [...] speed [-m] max_auto_speed`

`portcfgspeed -h`

Description Use this command to set the speed on a specified port or port range. This command disables and then re-enables the ports, and the ports come online with the new speed setting. The configuration is saved in nonvolatile memory and is persistent across switch reboots or power cycles. Use the **portShow** command to display supported port speed levels. Use the **portCfgShow** command to display configured speed settings.

On 16 Gbps-capable and 32 Gbps-capable platforms, the **portCfgSpeed** command checks if the requested speed is allowed, based on the combination configured for the octet that contains the port. If the speed is not supported by the current octet speed combination, this command exits with a warning message that displays the combination compatible with desired speed. Use the **portCfgOctetSpeedCombo** command to set the suggested combination before re-executing the **portCfgSpeed** command.

You can identify a single port to be configured by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers (decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

Notes This configuration cannot be set on VE/VEX_Ports. For a virtual FC port, the speed is always 10 GbE and port speed autonegotiation is not possible.

Speed configuration is not applicable to FCoE ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot For bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port Configures a single port identified by the port number relative to the slot on bladed systems.

portCfgSpeed

| | |
|---|--|
| -i <i>index1</i> [- <i>index2</i>] | Configures a port or a range of ports identified by port index numbers. Multiple port ranges are not supported with this command. |
| -f | Ignores nonexistent ports. This operand is valid only with the -i option. |
| -x <i>hex1</i> [- <i>hex2</i>] | Configures a port or a range of ports identified by port numbers, index number in hexadecimal format. You may specify multiple port ranges separated by a space, for example, -x 21-26 28-3c. |
| -slot <i>slot1</i> [- <i>slot2</i>] | Configures all ports on a slot or on a range of slots, for example, -s 3-5. You may specify multiple slot ranges separated by a space, for example, -s 3-5 8-10. |
| speed | Sets speed for the specified ports. This operand is required. Valid values are one of the following. |
| 0 | Autosensing mode (hardware). The port automatically configures for maximum speed. |
| ax | Autosensing mode (hardware). The port automatically configures for maximum speed with enhanced retries (not supported on 16G and 32G platforms). |
| s | Auto-sensing mode (software). The port automatically configures for maximum speed with enhanced retries (not supported on 32G platforms). |
| 1 | The port is set at a fixed speed of 1 Gbps (not supported on 10G, 16G, and 32G platforms). |
| 2 | The port is set at a fixed speed of 2 Gbps (not supported on 32G platforms). |
| 4 | The port is set at a fixed speed of 4 Gbps. |
| 8 | The port is set at a fixed speed of 8 Gbps. |
| 10 | The port is set at a fixed speed of 10 Gbps (not supported on 8G platforms). |
| 16 | The port is set at a fixed speed of 16 Gbps (not supported on 8G platforms). |
| 32 | The port is set at a fixed speed of 32 Gbps (not supported on 8G and 16G platforms). |

| | |
|-----------|---|
| -m | Sets auto-negotiation maximum speed. This operand is optional. Valid values are one of the following. |
| 2 | Sets the maximum auto-negotiation speed to 2 Gbps (not supported on 32G platforms). |
| 4 | Sets the maximum auto-negotiation speed to 4 Gbps. |
| 8 | Sets the maximum auto-negotiation speed to 8 Gbps. |
| 16 | Sets the maximum auto-negotiation speed to 16 Gbps. |
| 32 | Sets the maximum auto-negotiation speed to 32 Gbps. |
| -h | Displays the command usage. |

Examples To set the speed of a port to 10 Gbps:

```
switch:admin> portcfgspeed 2/3 10
```

To set the speed of a port using the port index:

```
switch:admin> portcfgspeed -i 78 16
```

To set the speed of a port range using the port index:

```
switch:admin> portcfgspeed -i 24-38 8
```

To set all ports on slots 2 and 3 to 8 Gbps:

```
switch:admin> portcfgspeed -s 2-3 8
```

To set all ports on slots 2-3 and 9-12 to 10 Gbps:

```
switch:admin> portcfgspeed -s 2-3 9-12 10
```

To set the speed of all ports in a range:

```
switch:admin> portcfgspeed -x 1d-1e 8
```

To set the maximum auto-negotiation speed:

```
switch:admin> portcfgspeed 17 0 -m 4
```

SEE ALSO [portCfgOctetSpeedCombo](#), [portCfgShow](#), [portShow](#), [switchCfgSpeed](#), [portSwapDisable](#), [switchCfgSpeed](#), [portSwapShow](#), [switchShow](#)

portCfgTdz

Configures a port for Target Driven Zoning (TDZ).

Synopsis `portcfgtdz --enable [slot/]port_list`

`portcfgtdz --disable [slot/]port_list`

`portcfgtdz --show [slot/]port_list`

`portcfgtdz --help`

Description

Use this command to configure a port to support Target Driven Zoning. This command enables or disables the Target Driven Zoning configuration on the specified port and saves the configuration persistently. Target Driven Zoning can be configured on E_Ports, F_Ports, and L_Ports that are online or offline and it does not toggle the port to apply the configuration. After Target Driven Zoning is configured on the specified port, it allows the connected target device to configure Target Driven Peer Zones to be enabled and committed to the switch fabric.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port on which the Target Driven Zoning configuration must be enabled or disabled, followed by a slash (/).

port_list

Specifies one or more ports, relative to the slot on bladed systems. Use **switchShow** for a listing of valid ports. A port list should be enclosed in double quotation marks and can consist of the following:

- A single port, for example, "8" or "5/8" on blades systems.
- A port range where beginning and end port are separated by a dash, for example, "8-13" or "5/8-13" on blades systems. A port range cannot span multiple slots.
- A set of ports separated by commas, for example "3,5,7,8" or "5/3,5,7,8" on bladed systems.
- A wildcard * indicates all ports. The wildcard can be represented as "*" or "***".

--enable

Enables the Target Driven Zoning configuration on the specified port.

--disable

Disables the Target Driven Zoning configuration on the specified port.

--show

Displays the current Target Driven Zoning configuration status (ON or OFF) for the specified port.

--help

Displays the command usage.

Examples To enable Target Driven Zoning configuration on a port:

```
switch:admin> portcfgtdz --enable 8
```

To display Target Driven Zoning configuration status for a specific port:

```
switch:admin> portcfgtdz --show 8
Port    Mode
=====
8       ON
```

To enable Target Driven Zoning configuration on a range of ports:

```
switch:admin> portcfgtdz --enable 8-18
```

To enable Target Driven Zoning configuration on all ports:

```
switch:admin> portcfgtdz --enable ""
```

To disable Target Driven Zoning configuration on a port:

```
switch:admin> portcfgtdz --disable 8
```

SEE ALSO [portCfgShow](#), [switchShow](#), [zoneAdd](#), [zoneCreate](#), [zoneDelete](#), [zoneRemove](#), [zoneShow](#)

portCfgTrunkPort

Enables or disables trunking on a port.

Synopsis `portcfgtrunkport [slot]/port[,] mode`

Description Use this command to enable or disable trunking on a port. Use **switchCfgTrunk** to enable or disable trunking on all ports of a switch.

When the command is executed to update the trunking configuration, the port to which the configuration applies is disabled and subsequently re-enabled with the new trunking configuration. Traffic through the ports may be temporarily disrupted.

Disabling trunking fails if a Trunk Area (TA) is enabled on the port. Use the **portTrunkArea** command to remove the TA before disabling trunking.

Notes Enabling trunking requires an ISL Trunking license. You may disable trunking without a license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The Fabric OS port configuration commands are not supported on FCoE ports.

Virtual FC Ports do not support frame-based trunking. If trunking is enabled for one of these ports in a slot, and a Brocade FX8-24 Extension blade is inserted into the slot, the configuration is not honored and the system logs a RASLOG error. An attempt to enable trunking for a Virtual FC port in a slot that contains a Brocade FX8-24 is rejected.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/).

port

Specifies the port to be configured, relative to its slot for bladed systems. Use **switchShow** to display a listing of valid ports.

mode

Specify 1 to enable trunking on the specified port. Specify 0 to disable trunking on the port. This operand is required. Trunking is enabled by default, when a trunking license is present on the switch.

Examples To enable a port for trunking:

```
switch:admin> portcfgtrunkport 1/3, 1
```

SEE ALSO [portCfgShow](#), [portShow](#), [switchCfgTrunk](#), [portTrunkArea](#), [switchShow](#)

portCfgVEXPort

Configures a port as a VEX_Port connected to an FC-IP and sets and displays VEX_Port configuration parameters.

Synopsis `portcfgvexport [slot/]port [-port]`

`portcfgvexport [-a admin]`

`portcfgvexport [-f fabricid]`

`portcfgvexport [-r ratov]`

`portcfgvexport [-e edtov]`

`portcfgvexport [-d domainid]`

`portcfgvexport [-p pidformat]`

`portcfgvexport [-t fabric_parameter]`

`portcfgvexport [-m portmode]`

Description Use this command to configure a port as a VEX_Port, to display the port's VEX_Port configuration, or to change the configuration. If no optional parameter is specified, the command displays the currently configured values; otherwise, it sets the specified attribute to its new value. The port must be disabled prior to setting VEX_Port attributes. The port must be enabled before the port can become active following VEX_Port parameter changes. Use **portDisable** and **portEnable** to disable or enable the port.

When the port is not active, the preferred domain ID is configurable. The preferred domain ID is used by the VEX_Port's front phantom domain to request a domain ID from the principal switch. The domain ID received becomes the subsequent preferred domain ID, which is persistent and is displayed.

Notes In FOS v7.1.0 and later, both Open(IM1) and McDATA(IM2 and IM3) interop modes are not supported.

The fabric ID must be the same for every router port connected to the same edge fabric, and different for every edge fabric. If two ports are connected to the same fabric but have been assigned different fabric IDs, one of them will be disabled due to a fabric ID oversubscription. If two fabrics have been assigned the same fabric ID, one of them will be disabled due to a fabric ID conflict.

The front domain WWN field displays the WWN of the front domain. If the port is enabled and the state is "OK", the edge fabric principal switch domain ID and WWN also are displayed.

If the Fabric Parameter value is "Auto Negotiate", the port ID format, R_A_TOV, and E_D_TOV values display the negotiated values indicated by "(N)" next to them. The negotiated values are what the edge switch specifies in the ELP request. If the state is "Not OK", R_A_TOV and E_D_TOV display "Not Applicable". By default, all VEX_Ports are auto-ELP enabled.

If the Fabric Parameter attribute value is "User configured", port ID format, R_A_TOV, and E_D_TOV display the configured values.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

| | |
|----------------------------|---|
| <i>slot</i> | On bladed systems only, specifies the slot number followed by a slash (/). |
| <i>port[-port]</i> | Specifies a port or a port range, relative to the slot number on bladed systems, for example, 5/17-29. Use switchShow for a list of valid ports. When executed with [<i>slot</i>] <i>port</i> [- <i>port</i>] only, the command displays the current VEX_Port configuration. |
| -a admin | Enables or disables the specified port as a VEX_Port. Valid values are 1 (enable as VEX_Port), 2 (disable as VEX_Port and enable as non-VEX_Port). portCfgDefault may also be used to disable VEX_Ports. |
| -f fabricid | Specifies the fabric ID. Valid values are 1-128. |
| -r ratov | Specifies the R_A_TOV used for port negotiation. Valid values are 2000 to 120000. This operand is only applicable if the "Fabric Parameter" attribute value is not "Auto Negotiate". |
| -e edtov | Specifies the E_D_TOV used for port negotiation. Valid values are 1000 to 60000. This operand is only applicable if the "Fabric Parameter" attribute value is not "Auto Negotiate". |
| -d domainid | Specify the preferred domain ID. For Brocade native mode (-m 0), Valid values are 1 to 239. |
| -p pidformat | Specifies the Port ID format. Valid values are 0-native, 1-core, 2-extended edge. This operand is applicable only when <i>port mode</i> is set to 0 (native mode). If <i>port mode</i> is not "Brocade Native", the Port ID format displays as "Not applicable". |
| -t fabric_parameter | Enables or disables negotiation of the fabric parameters. Valid values are: 1-enable, 2-disable. |
| -m port mode | Specifies the Port mode. The -m option enforces the same port mode for all the ports connected to the same edge fabric. If the -m option is selected, the port mode is compared to the online ports. If the modes are different, an error message is displayed, and the command fails. Valid values are as follows: |
| 0 | Brocade Native mode. |

Examples To display the VEX_Port configuration of port 2/16:

```
switch:admin> portcfgvlexport 2/16

      Port   2/16   info
Admin:                enabled
State:                OK
Pid format:          core(N)
Edge Fabric ID:      16
Front Domain ID:     160
Front WWN:           50:06:06:9e:20:9f:ce:10
Principal Switch:    7
principal WWN:       10:00:00:60:69:c0:05:8a
Fabric Parameters:   Auto Negotiate
R_A_TOV:             9000(N)
E_D_TOV:             2000(N)
Authentication Type: DHCHAP
DH Group: 4
Hash Algorithm: SHA-1
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

To set the fabric ID of port 2/21 to 5 and the port ID format to core:

```
switch:admin> portcfgvlexport 2/21 -f 5 -p 1
```

To configure port 2/20 as a VEX_Port and set the fabric ID to 4:

```
switch:admin> portcfgvlexport 2/20 -a 1 -f 4
```

To disable fabric parameter negotiation on port 2/20 of a VEX_Port:

```
switch:admin> portcfgvlexport 2/20 -t 2
```

SEE ALSO [portCfgEXPort](#), [portDisable](#), [portEnable](#), [portShow](#)

portChannelShow

Displays the port channels from a domain to the remote domains.

Synopsis `portchannelshow [domain]`

`portchannelshow --help`

Description Use this command to display information about port channels from a domain to one or all reachable remote domains. The command output includes the following information:

domain(s) in the fabric

Total number of domains in the fabric.

Local Domain ID

The domain number of local switch.

Domain

The domain number of destination switch.

Name

The name of the destination switch.

WWN

The world wide name (WWN) of the destination switch.

Port Channel

The list of port-channel interfaces. For each trunk, only the port number of the master port is displayed. This field displays "None" during the following conditions:

- Policy-Based Routing (PBR) is enabled on the switch.
- Traffic Isolation (TI) zone is applied on the egress port.
- Port channel is used by 8 Gbps-capable ingress ports.
- Port channel contains ICL ports or it is used by other ICL ingress ports.
- Port channel is not used by any ingress port.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

domain

Specifies the remote domain to display port channels between the switch in which command is executed to the specified remote domain. You can specify *domain* in one of the following formats:

domain_id

Specifies the domain number of destination switch.

name

Specifies the name of destination switch.

wwn

Specifies the WWN of destination switch.

--help

Displays the command usage.

Examples To display port channels for all reachable domains:

```

switch:admin> portchannelshow
6 domain(s) in the fabric; Local Domain ID: 4

Domain:      1
Name:        sw0
WWN:         10:00:00:05:33:c1:26:00
Port Channel: None

Domain:      2
Name:        DCX_35_F_128
WWN:         10:00:00:05:1e:38:e5:23
Port Channel:
Ports:       384, 385, 386, 387, 400, 401, 402, 403,
             417, 418, 419, 432, 433, 434, 435

Domain:      3
Name:        SW_122_F_128
WWN:         10:00:00:05:1e:9b:10:5b
Port Channel:
Ports:       111, 248

Domain:      5
Name:        SW_65_F128
WWN:         10:00:00:05:1e:5c:f6:fd
Port Channel:
Ports:       384, 385, 386, 387, 400, 401, 402, 403,
             417, 418, 419, 432, 433, 434, 435

Domain:      6
Name:        SW_121_F_128
WWN:         10:00:00:05:1e:9c:32:cc
Port Channel: None

```

To display port channels for a specific remote domain:

```

switch:admin> portchannelshow 10:00:00:05:1e:38:e5:23
Domain:      2
Name:        DCX_35_F_128
WWN:         10:00:00:05:1e:38:e5:23
Port Channel:
Ports:       384, 385, 386, 387, 400, 401, 402, 403,
             417, 418, 419, 432, 433, 434, 435

```

SEE ALSO [islShow](#), [trunkShow](#)

portCmd

Diagnoses intelligent ports.

Synopsis

```
portcmd --ping [slot][slot]ge port
  -s src_ip -d dst_ip [-x | -crossport]
  [-n num_requests][-q service_type] [-t ttl][-w wait_time]
  [-z size][-v vlan_id [-c L2 Class-of-Service]]

portcmd --traceroute [slot][slot]ge port
  -s src_ip -d dst_ip [-x | -crossport]
  [-m max_hops][-f first_ttl][-q type_of_service][-w timeout]
  [-z size] [-v vlan_id [-c L2 Class-of-Service]]

portcmd --tperf [slot]veport -sink | -source
  [-high | -medium | -low][-time duration] [-unidirectional]
  [-random] [-pattern pattern] [-size pdu_size]
  [-interval interval]

portcmd --pmtu [slot] ge port -s src_ip -d dst_ip
  -q DSCP -c L2 Class-of-Service -v -y

portcmd --wtool wt-id | all action port arguments

portcmd --help
```

Description

Use this command to ping or trace a route to a destination IP host from an intelligent GbE port, or to determine the path characteristics between a local data source and a remote data sink.

When issued with the **--tperf** option, this command determines the path characteristics to a remote host or tunnel destination. TPerf generates statistics every 30 seconds by default unless you specify a different value for **-interval**. The output displays the following information:

| | |
|---------------------------|--|
| Tunnel ID | Numeric identifier for the TPerf tunnel. |
| Traffic Priority | High, Medium, or Low. |
| bytes tx | Number of bytes transmitted. |
| bytes rx | Number of bytes received. |
| PDU tx | Number of protocol data units transmitted. |
| PDU rx | Number of protocol data units received. |
| bad CRC headers rx | Number of bad CRC headers received. |

| | |
|----------------------------|--|
| bad CRC payloads rx | Number of bad CRC payloads received. |
| out of seq PDUs rx | Number of out-of-sequence PDUs received. |
| flow control count | Flow control count. |
| last rtt | Last Round trip in milliseconds (RT traffic only). |

VLAN tagging ensures that test traffic traverses the same path as real FCIP traffic. A VLAN tag entry must exist prior to issuing the **--ping** or **--traceroute** commands; this includes both the local and remote sides.

Notes End-to-end path characterization is not supported if an IPSec-enabled tunnel exists that uses the same source/local IP address.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------------|---|
| slot | For bladed systems only, specifies the slot number of the port to be configured, followed by a slash (/). |
| ge port | Specifies the port number of the GbE port on the blade. For the Brocade 7840 switch and the Brocade SX6 extension blade, specify GbE port number along with the DP number, for example, portcfg ipif ge0.dp0 . The valid DP numbers are dp0 and dp1. |
| --ping | Pings a destination IP address from one of the source IP interfaces on the GbE port. Valid arguments and their values include the following: |
| -s src_ip | Specifies the source IP address that originates the ping request. IPv6 addresses are supported. |
| -d dst_ip | Specifies the destination IP address to which to target the ping request. IPv6 addresses are supported. |
| -x -crossport | Specifies the source IP address as a crossport. A crossport is the nonlocal (secondary) XGE port to a VE_Port group that provides failover capabilities for the tunnel configured on the local XGE port. This operand is optional and supported only on the Brocade FX8-24 XGE ports. |
| -n num_requests | Specifies the number of ping requests. Valid values are 1 to 255. The default is 4. This operand is optional. |

- q *service_type***
Specifies the type of service in the ping request. The default is 0 and *service_type* can be hex (0x0-0xFF) or decimal (0-255). This operand is optional.
- t *tll***
Specifies the time to live. Valid values are 1 to 255. The default is 100. This operand is optional.
- w *wait_time***
Specifies the time to wait for the response of each ping request in seconds. The default is 5 seconds and the maximum wait time is 9 seconds. This operand is optional.
- z *size***
Specifies the default packet size to a fixed size in bytes. The default is 64 bytes. In an IPv4 environment, the ICMP/IP header occupies 28 bytes. In an IPv6 environment it occupies 48 bytes. The total size, including ICMP/IP headers (28 or 48 bytes without IP options) cannot be greater than the IP MTU configured on the interface. This operand is optional.
- v *vlan_id***
Specifies the VLAN ID. Values must be in the range between 1 and 4094. There is no default value. Note that a VLAN tag entry must exist on the local and remote sides prior to issuing the **-v** option. A VLAN Tag table entry is dynamically maintained by the **ipperf** application. See the **portCfg** help page for details on creating a VLAN tag table. This option is supported on Brocade FX8-24 blade only.
- c *L2 Class-of-Service***
Specifies Class of Service/Priority, as defined by IEEE 802.1p. Values must be in the range between 0 and 7. The default is 0. This operand is optional with the **-v** option.
- traceroute**
Traces the IP router hops used to reach the host *dst_ip* from one of the source IP interfaces on the GbE port. Valid arguments include the following:
- s *src_ip***
Specifies the local IP address to use for sourcing the probe packets. IPv6 addresses are supported.
- d *dst_ip***
Specifies the destination IP address to which to probe the IP router path. IPv6 addresses are supported.
- x | -crossport**
Specifies the source IP address as a crossport. A crossport is the nonlocal (secondary) XGE port to a VE_Port group that provides failover capabilities for the tunnel configured on the local XGE port. This operand is optional and supported only on the Brocade FX8-24 XGE ports.
- m *max_hops***
Specifies the maximum hop limit used in the outgoing probe packets. The default value is 30 hops. This operand is optional.

- f *first_ttl***
Specifies the starting point for the "time to live" parameter. The default is 1. The command skips processing for those intermediate gateways that are less than the *first_ttl* hops. This operand is optional.
- q *service_type***
Specifies the type of service in the **traceroute** request. The default is 0 and *service_type* must be an integer from 0 to 255. This operand is optional.
- w *wait_time***
Sets the time, in milliseconds, to wait for a response to a probe. The default is 5000 milliseconds. The maximum wait time is 9000 milliseconds (9 seconds). This operand is optional.
- z *size***
Specifies the size, in bytes, of the trace route packet to use. The default is 64 bytes. In an IPv4 environment, the ICMP/IP header occupies 28 bytes. In an IPv6 environment, it occupies 48 bytes. The total size, including ICMP/IP headers (28 or 48 bytes without IP options), cannot be greater than the IP MTU configured on the interface. This operand is optional.
- v *vlan_id***
Specifies the VLAN ID. Values must be in the range between 1 and 4094. There is no default value. Note that a VLAN tag entry must exist on the local and remote sides prior to issuing the **-v** option. See the **portCfg** help page for details on creating a VLAN tag table. This option is supported on Brocade FX8-24 blade only.
- c *L2 Class-of-Service***
Specifies Class of Service/Priority, as defined by the IEEE 802.1p specification. Values must be in the range between 0 and 7. The default value is 0. This operand is optional with the **-v** option.
- help**
Displays command usage.
- tperf**
Determines the path characteristics to a remote host or tunnel destination. The **-tperf** option requires two separate Brocade FX8-24 blades to function. One blade plays the role of a data sink and the other blade plays the role of the data source. TPerf also requires that you define a tunnel as a TPerf tunnel. Refer to **portCfg fcipunnel** for more information. This option is supported on Brocade FX8-24 blade only.
- sink | -source**
Designates the switch to function either as a data sink or a data source. This operand is required.
- When **-sink** is specified, TPerf begins to respond to traffic sent by the switch acting as the data source. The process continues to run until it is either terminated by user intervention (**Ctrl +C**) or, if a value for duration is specified with the **-t** option, until the process completes the set time frame.

The following optional arguments are ignored on the data sink, because it services all requests from the data source: **-high**, **-medium**, **-low**, **-unidirectional**, **-random**, **-pattern**, and **-size**.

When **-source** is specified, TPerf generates traffic until it is interrupted by user intervention (**Ctrl + c**) or, if a value for duration is specified with the **-t** option, until the process completes the set time frame. The TPerf module on the remote host will immediately begin generating traffic; it is therefore imperative that the data sink has been started on the opposing switch before the data source is started on the local switch.

The following arguments are optional:

- high**
Generates high priority traffic.
- medium**
Generates medium priority traffic.
- low**
Generates low priority traffic.

If no traffic priority is specified only medium priority traffic is generated.
- time *duration***
Specifies the duration of the TPerf traffic flow in seconds. If a value for duration is not specified, the process continues to run until it is terminated with **Ctrl + C**.
- unidirectional**
Generates traffic in one direction only. The default is round-trip.
- random**
Specifies a random protocol data unit (PDU) size between 2048 and the size of the send request. Refer to **-size** below.
- pattern *pattern***
Specifies the test data pattern for the payload as one of the following values:
 - 0 | notspecified**
No pattern is specified. TPerf applies whatever is already set or in memory. This is the default value.
 - 1 | allzeros**
The specified pattern is "all zeros".
 - 2 | allones**
The specified pattern is "all ones".
 - 3 | incrementingbyte**
The specified pattern is "incrementing byte".
 - 4 | random**
The specified pattern is "random".
 - 5 | jitter**
The specified pattern is "jitter".

- size pdu_size**
Specifies the PDU size to use (not including headers). The valid range is between 2048 and 10112 bytes. The default is equivalent to the maximum segment size (MSS). This is the maximum size if the **-random** option is specified.
- interval interval**
Specifies the interval at which the statistics display is refreshed, in seconds. The default is 30 seconds.
- pmtu**
Generates the ICMP messages to discover the maximum MTU available on a specific network path. It is represented as the maximum number of bytes that can traverse the network path. This option is supported only on Brocade 7840 switch and the Brocade SX6 extension blade. Valid arguments include the following:
- s src_ip**
Specifies the local IP address to use for sourcing the probe packets. IPv6 addresses are supported.
- d dst_ip**
Specifies the destination IP address to which to probe the IP router path. IPv6 addresses are supported.
- c L2 Class-of-Service**
Specifies Class of Service/Priority, as defined by IEEE 802.1p. Values must be in the range between 0 and 7. The default is 0.
- q DSCP**
Specifies the DSCP marking. Values must be in the range between 0 and 63.
- v**
Displays verbose output.
- y**
Displays the short output with summary lines at the beginning and end.
- wtool**
Generates traffic over a pair of IP addresses to test the link for any network issues. Determines the health of a network link before deploying it for use in a circuit. This option is supported only on Brocade 7840 switch and the Brocade SX6 extension blade.
- The syntax for **portcmd --wtool** is as follows:
- portcmd --wtool wt-id | all action port arguments**
- This command supports following actions and arguments:
- create**
Creates a WAN test session. This command supports following options:
- a | --admin-status enable | disable**
Enables or disables the WAN test session.
- s | --src src_ip**
Specifies the local IP address to use for sourcing the probe packets. IPv6 addresses are supported.

- d | --dst *dst_ip***
Specifies the destination IP address to which to probe the IP router path. IPv6 addresses are supported.
- r | --rate *rate***
Specifies the test rate in kbps.
- t | --time *time***
Specifies the test time in seconds.
- i | --ipsec-policy *ip-sec_policy***
Specifies the policy name for an ipsec configuration.
- b | --bi-directional**
Specifies the test session is running on bi-directional traffic. By default, the WAN test session is configured for uni-directional traffic.
- u | --uni-directional**
Specifies the test session is running on uni-directional traffic.
- dscp *dscp***
Sets Differentiated Services Code Point (DSCP) markings for the WAN test session.
- l2cos *l2cos***
Sets the Layer 2 Class of Service (L2CoS) value for the WAN test session.

modify

Modifies an established WAN test session. This command supports following options:

- r *rate***
Specifies the test rate in kbps.
- t *time***
Specifies the test time in seconds.
- b**
Specifies the test session is running on bi-directional traffic. By default, the WAN test session is configured for uni-directional traffic.
- c**
Clears the statistics associated with the WAN test.

show

Displays the WAN test session statistics. This command supports following options:

- summary**
Displays statistics for active WAN test sessions.
- connection**
Displays statistics for established WAN test sessions.
- detail**
Displays detailed statistics for active WAN test sessions.
- historical**
Displays statistics for past WAN test sessions.

| | |
|-----------------|---|
| --all-ls | Displays statistics for all configured WAN test sessions from all logical switches. |
| start | Starts the WAN test session associated with the test ID. This command supports following option: |
| -t | Specifies the test time in seconds. Use this option if the test time is not configured with create or modify options. |
| stop | Stops the WAN test session associated with the test ID. Use all to stop all WAN test sessions. |
| delete | Deletes the WAN test session associated with the test ID. Use all to delete all WAN test sessions. |
| --help | Displays command usage. |

Examples To verify if packets can be sent to the destination IP address with maximum *wait_time* specified:

```
switch:admin> portcmd --ping 12/ge0 \  
-s 2007:7:30:32:227:138:10:120 \  
-d 2007:7:30:32:227:77:0:60 -w 29000  
Pinging 2007:7:30:32:227:77:0:60 from ip interface \  
2007:7:30:32:227:138:10:120 on 12/ge0 with 64  
  
bytes of data  
Reply from 2007:7:30:32:227:77:0:60: bytes=64 rtt=0ms ttl=255  
Reply from 2007:7:30:32:227:77:0:60: bytes=64 rtt=1ms ttl=255  
Reply from 2007:7:30:32:227:77:0:60: bytes=64 rtt=0ms ttl=255  
Reply from 2007:7:30:32:227:77:0:60: bytes=64 rtt=0ms ttl=255  
  
Ping Statistics for 2007:7:30:32:227:77:0:60:  
Packets: Sent = 4, Received = 4, Loss = 0 ( 0 percent loss)  
Min RTT = 0ms, Max RTT = 1ms Average = 0ms
```

To trace the IP router hops used to reach the remote (with packet size specified):

```
switch:admin> portcmd --traceroute 12/ge0 \  
-s 2007:7:30:32:227:138:10:120 \  
-d 2007:7:30:32:227:77:0:60 -z 1452  
Traceroute to 2007:7:30:32:227:77:0:60 from IP interface  
2007:7:30:32:227:138:10:120 on 12/0, 30  
  
hops max  
1 1 ms 0 ms 0 ms  
Traceroute complete.
```

To verify if packets can be sent to the destination IP address using VLAN tagging with the **-c** option.

```
switch:admin> portcmd --ping 8/ge0 \  
-s 192.168.10.1 -d 192.168.20.1 -v 10 -c 3
```

To trace the IP router hops used to reach the remote host using VLAN tagging with the `-c` option.

```
switch:admin> portcmd --traceroute 8/ge0 \
-s 192.168.10.1 -d 192.168.20.1 -v 10
```

To ping a failover crossport:

```
switch:admin> portcmd --ping 8/xge0 \
192.168.11.20 -d 1.1.1.1 -x
```

To trace the route to a failover crossport:

```
switch:admin> portcmd --traceroute 8/xge0 \
192.168.11.20 -d 1.1.1.1 -x
```

To set the path characteristic to source mode on the remote host using VLAN tagging with the `-c` option.

```
switch:admin> portcmd --iperf 8/ge0 \
-s 192.168.10.1 -d 192.168.20.1 -S -v 10 -c 3
```

To create a TPerf data sink on VE_Port 16:

```
switch:admin> portcmd --tperf 16 -sink

TPerf has been configured successfully for 16
TPerf is servicing requests on 16 priority: high
TPerf is servicing requests on 16 priority: medium
TPerf is servicing requests on 16 priority: low
Tperf data source can now be started
```

To configure a Brocade FX8-24 blade as a TPerf data source (this example generates round-trip high, medium, and low priority traffic with a fixed PDU size of 2084 bytes. The payload pattern is a jitter pattern and the checksum is checked by the sink):

```
switch:admin> portcmd --tperf 16 -source \
-size 2048 -pattern 5

TPerf has been configured successfully for 16
TPerf is generating traffic on 16 priority: medium
*****
Tunnel ID: 16

High Priority Medium Priority Low Priority
bytes tx          0          5280520          0
bytes rx          0          407240           0
PDUs tx           0          10183            0
PDUs rx           0          10182            0
out of seq PDUs rx 0           0                0
flow control count 0           0                0
last rtt          0           5                0
*****
```

To configure a TPerf source so it generates round-trip high and low priority traffic with a random PDU size from 2048 to 10112 bytes (1 MSS) and with a payload pattern determined by previous settings saved in memory:

```
switch: admin> portcmd --tperf 16 -source -high -low -random
TPerf has been configured successfully for 16
TPerf is generating traffic on 16 priority: high
TPerf is generating traffic on 16 priority: low
*****
Tunnel ID: 16
```

| | High Priority | Medium Priority | Low Priority |
|--------------------|---------------|-----------------|--------------|
| bytes tx | 19800708 | 0 | 19168848 |
| bytes rx | 187200 | 0 | 183640 |
| PDUs tx | 4686 | 0 | 4598 |
| PDUs rx | 4685 | 0 | 4597 |
| out of seq PDUs rx | 0 | 0 | 0 |
| flow control count | 0 | 0 | 0 |
| last rtt | 5 | 0 | 5 |

```
*****
```

To configure a TPerf source to generate one-way low priority traffic with a random PDU size from 2048 to 4000 bytes and a payload pattern of ones.

```
switch:admin> portcmd --tperf 16 -source -low \
-size 4000 -random -unidirectional -pattern 2

TPerf has been configured successfully for 16
TPerf is generating traffic on 16 priority: low
*****
Tunnel ID: 16
```

| | High Priority | Medium Priority | Low Priority |
|--------------------|---------------|-----------------|--------------|
| bytes tx | 0 | 0 | 19800708 |
| bytes rx | 0 | 0 | 0 |
| PDUs tx | 0 | 0 | 1434686 |
| PDUs rx | 0 | 0 | 0 |
| out of seq PDUs rx | 0 | 0 | 0 |
| flow control count | 0 | 0 | 0 |
| last rtt | N/A | N/A | N/A |

```
*****
```

To specify test data patterns for the source payload:

```
switch:admin> portcmd --tperf 16 -source -pattern random
switch:admin> portcmd --tperf 16 -source -pattern 4

switch:admin> portcmd --tperf 16 -source -pattern jitter
switch:admin> portcmd --tperf 16 -source -pattern 5
```

To discover the maximum MTU available on a switch:

```
switch:admin> portcmd --pmtu ge3.dp0 -s 192.168.0.30 -d 192.168.0.130
Path MTU detection from 192.168.0.30 to 192.168.0.130
Note: a response can take up to 32 seconds

Path MTU returned:1500
```

To create a WAN test session:

```
switch:admin> portcmd --wtool 0 create -d 76.196.3.76 -s 77.195.3.77 -r 5000000
Operation Succeeded
```

To start a WAN test session:

```
switch:admin> portcmd --wtool 0 start -t 1200
Operation Succeeded
```

To display WAN test session for a specified session:

```
switch:admin> portcmd --wtool 0 show
wantool-id: (0)
=====
State           : Established
Up Time        : 1m6s
Run Time       : 0s
Time remaining  : 0s
IP Addr (L/R)  : 76.196.3.76 - 77.195.3.77
PMTUD          : Disabled
Comm Rate      : 5000000 Kbps (610.35 MB/s)
Tx rate        : 114.06 Kbps (0.01 MB/s)
Rx rate        : 102.66 Kbps (0.01 MB/s)
Tx Utilization : 0.00%
Rx Utilization : 0.00%
RTT (Min/Max)  : 0.10ms/0.28ms
RTT VAR (Min/Max) : 0.09ms/0.28ms
Local Session Statistics
Tx pkts        : 0
Peer Session Statistics
Rx pkts        : 0
Ooo pkts       : 0
Drop pkts      : 0 (0.00%)
```

To display active WAN test sessions:

```
switch:admin> portcmd --wtool show --d
WTool Session: 24.0 (DP0)
=====
Admin / Oper State : Enabled / Running
Up Time           : 1m24s
Run Time          : 1m23s
Time Out          : 3m50s
Time Remaining    : 13m37s
IP Addr (L/R)    : 192.168.3.20 ge3 <-> 192.168.3.10
IP-Sec Policy     : (none)
PMTU Discovery (MTU) : disabled (1500)
Bi-Directional   : disabled
L2CoS / DSCP     : (none) / (none)
Configured Comm Rate : 5000000 kbps
Peer Comm Rate    : 5000000 kbps
Actual Comm Rate   : 5000000 kbps
Tx rate           : 5000000.00 Kbps ( 610.00 MB/s)
Rx rate           : 5000000.00 Kbps ( 610.00 MB/s)
Tx Utilization    : 100.00%
Rx Utilization    : 100.00%
RTT (Min/Max)    : 1 ms/1 ms
```



```
RTT VAR (Min/Max)      : 1 ms/1 ms
Local Session Statistics
Tx pkts                : 810024
Peer Session Statistics
Rx pkts                : 792029
Ooo pkts               : 0
Drop pkts              : 0 (0.00%)
(output truncated)
```

To modify a WAN test session:

```
switch:admin> portcmd --wtool 0 modify --bi
Operation Succeeded
```

To stop a WAN test session:

```
switch:admin> portcmd --wtool 0 stop
Operation Succeeded
```

SEE ALSO [portCfg](#), [portShow](#)

portDebug

portDebug

Sets debug level and verbose level of port modules.

Synopsis `portdebug dbg_lvl, vbs_lvl`

Description Use this command to set the debug level and verbose level of port modules.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

dbg_lvl Specify the debug level to be set for port modules; valid values are 1 to 5.

vbs_lvl Specify the verbose level to be set for port modules; valid values are 1 to 5.

Examples To set debug level and verbose level of port modules:

```
switch:admin> portdebug 3 4
```

SEE ALSO [dbgShow](#)

portDecom

Decommissions and disables an E_Port without frame loss.

Synopsis

portdecom

portdecom [*slot*]/*port*

portdecom --help

Description

Use this command to disable an E_Port without frame loss. The E_Port can either be a single redundant E_Port or an E_Port which is part of a trunk group. The port must be an operational E_Port with at least one other redundant E_Port available to reach all domains that the E_Port can reach. Or the port must be an operational E_Port which is part of a trunk group with at least one other member in the trunk. The last member of the trunk group may also be decommissioned, if there are other available redundant E_Ports.

Upon successful completion of the command all traffic flows on the port will have been moved to other redundant paths and the port will be placed in a persistently disabled state. To bring the port back online use the **portCfgPersistentEnable** command.

Notes

The execution of this command requires that the lossless feature is enabled on both the local switch and the remote switch on the other end of the E_Port. This command will fail for online ports, if the port is not a fully operational E_Port or trunk port, lossless is not enabled on the local or remote switch, or no alternate redundant path is available to route the traffic. A redundant path must be an E_Port connection between the same pair of switches as the E_Port that is being decommissioned. Use the **topologyShow** command to verify that alternate paths are available to reach all domains.

When decommissioning the member of a trunk group, all members must have an equal link cost. If a member does not have an equal cost, the cost of the trunk as a whole may change after decommissioning the port and cause routes to be changed. Refer to the **linkCost** command for more information.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the ports to be disabled, followed by a slash (/).

port

Specifies the E_Port to be decommissioned.

Examples

To decommission a port

```
switch:user> portdecom 2/5
```

```
The port has been decommissioned and persistently disabled.
```

SEE ALSO

[dlsSet](#), [dlsShow](#), [linkCost](#), [portCfgPersistentEnable](#), [topologyShow](#)

portDisable

Disables a port or a range of ports.

Synopsis

```
portdisable [slot]port1[-port2] [...]
```

```
portdisable -i [index1[-index2] [...]] [-f]
```

```
portdisable -x [hex1[-hex2] [...]]
```

```
portdisable -slot [slot1[-slot2] [...]]
```

```
portdisable [slot]port -r disable_reason_string
```

```
portdisable -h
```

Description Use this command to disable a port or a range of ports. If the port is connected to another switch when disabled, the fabric may reconfigure. Devices connected to this port can no longer communicate with the fabric.

You can identify a single port to be disabled by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers(decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

The front panel LED of a disabled port flashes amber in a two-second cycle. If the port was online before being disabled, a state transition will be indicated in one of the following ways: RSCN, SNMP trap, or Web pop-up window.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports. To disable an FCoE port, use **fcoe --disable**.

Operands This command has the following operands:

| | |
|-------------------------------------|---|
| <i>slot</i> | On bladed systems only, specifies the slot number of the ports to be disabled, followed by a slash (/). |
| <i>port1</i> [- <i>port2</i>] | Disables a single port or a range of ports identified by port numbers. The port range cannot span slots, but you can specify multiple port ranges pairs separated by a space, for example, 3/1-4 4/7-9. |
| -i <i>index1</i> [- <i>index2</i>] | Disables a port or a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, -i 33-38 40-60. |
| -f | Ignores nonexistent ports. This operand is valid only with the -i option. |

- x [hex1 [-hex2]]**
Disables a port or a range of ports identified by port numbers, index number in hexadecimal format. You may specify multiple port ranges separated by a space, for example, -x 21-26 28-3c.
- slot [slot1[-slot2]]**
Disables all ports on a slot or on a range of slots, for example, -s 3-5. You may specify multiple slot ranges separated by a space, for example, -s 3-5 8-10.
- r *disable_reason_string***
Specifies the reason for disabling the port. The string can be up to 16 characters long and must be enclosed in double quotation marks.
- h**
Displays the command usage.

Examples

To disable a single port.

```
switch:admin> portdisable 2/4
```

To disable a range of ports:

```
switch:admin> portdisable 2/4-8
```

To disable multiple port ranges:

```
switch:admin> portdisable 2/24-26 3/10-12 4/3-4
```

To disable a port by specifying its index number:

```
switch:admin> portdisable -i 176
```

To disable a range of ports by specifying the corresponding port index range:

```
switch:admin> portdisable -i 170-176
```

To disable multiple ports by specifying multiple port index ranges:

```
switch:admin> portdisable -i 30-36 170-176
```

To disable range of ports by specifying port index numbers in hexadecimal format:

```
switch:admin> portdisable -x 1d-1e
```

To disable all ports on slots 3-5:

```
switch:admin> portdisable -s 3-5
```

To disable all ports on slots 3-5 and 9-12:

```
switch:admin> portdisable -s 3-5 9-12
```

portDisable

To disable a port with a disable reason:

```
switch:admin> portdisable 2/2 -r "laser fault"
```

SEE ALSO

[portCfgPersistentDisable](#), [portCfgPersistentEnable](#), [portEnable](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portDPortTest

Initiates or terminates a D_Port test, or displays D_Port test results.

Synopsis

```

portdporttest --start | --setarg [-nframes number | -time time]
  [-framesize size] [-pattern pat_name |
  -payload payload_pattern] [-fec] [-cr] [slot]/port_list

portdporttest --stop [slot]/port_list

portdporttest --restart [slot]/port_list

portdporttest --set [ewrap | owrap] [slot]/port

portdporttest --clear [ewrap | owrap | all] [slot]/port

portdporttest --show [slot]/port_list | -detail [slot]/port_list | all

portdporttest --exit [slot]/port_list

portdporttest --help

```

Description Use this command to manually terminate, re-initiate test, or stop test on a diagnostic port (D_Port). The **portdporttest --show** command displays the result of the D_Port tests. The port must be configured as a D_Port and physically connected to a second D_Port on a remote switch. Refer to the **portCfgDPort** command for more information. The **portdporttest --start** command only needs to be issued on one end of the link.

The D_Port test performs the following diagnostics:

- An electrical loopback test (supported only on 16Gb and 32Gb SFPs capable of electrical loopback)
- An optical loopback test (supported only on 16Gb and 32Gb SFPs capable of optical loopback)
- A link traffic test
- A link distance measurement

Use the **--setarg** option to set the test parameters and exit. Use the **--restart** option to restart the test with the already configured parameters.

Use the **--show** option with a port operand to display the test status for the specified ports at any time during the test or after the diagnostics has completed. The command output includes the following information for each configured D_Port:

| | |
|---------------------|---|
| Slot | The slot number on the local side of the link. |
| Port | The port number on the local side of the link |
| Remote WWNN | The world wide node name of the switch on the remote side of the link. |
| Remote WWPNN | The world wide port name of the Host Bus Adapter (HBA) or device port on the remote side of the link. |

| | |
|------------------------------------|--|
| Remote port index | The port number on the remote side of the link |
| Mode | The testing mode: Automatic (initiated due to port online event) or Manual (manually restarted using this command). |
| No. of test frames | Number of frames sent for link traffic test. |
| Test frame size | The test frame size in bytes. |
| Pattern | The predefined pattern name. |
| Payload | The user-defined payload pattern. |
| FEC (enabled/option/active) | The forward error correction (FEC) status. This option is not supported for D_Port HBA. |
| CR (enabled/option/active) | The credit recovery (CR) status. This option is not supported for D_Port HBA. |
| Start time | Date and time of test start |
| End time | Date and time of test completion |
| Status | On the link initiator, this field displays the overall test status as PASSED, FAILED, SKIPPED, IN PROGRESS, NOT STARTED, or STOPPED. On the link responder, the status displays PASSED, FAILED, STOPPED, or RESPONDER. |
| Test | For each sub-test, the command displays the Start time, test result (PASSED, FAILED, IN PROGRESS, NOT STARTED, STOPPED, or RESPONDER), estimated time in seconds, and diagnostic comments. Comments include the following: <ul style="list-style-type: none"> <i>No SFP or chip support</i> The test result is skipped because the SFP or port is not capable of D_Port functionality. <i>See remote port results</i> The test result related to this comment is RESPONDER. The port is in responder state and test results will be available at the initiator port on the remote side of the link. <i>Remote port is not ready to start the test</i> The test failed because the remote port is not ready to start the D_Port test. <i>No remote SFP or chip support</i> The test is skipped because the D_Port test is run between a QSFP and SFP port connected through break-out cable or when the remote SFP is not capable of running the optical loopback test. |

Roundtrip link latency

Displays roundtrip link latency in nanoseconds or unknown when the test is in progress. Roundtrip link latency is the time required for the light to travel from a given port to the port at the remote side of the link and back again.

Estimated cable distance

The estimated physical distance of the interswitch link in meters. The Link distance for SPFs or ICL ports which do not support E_WRAP and O_WRAP have a precision of +/- 50 meters for 8 Gbps LWL SFPs, 8 Gbps ELWL SFPs, 10 Gbps SFPs, ICL ports, and QSPF ports and +/- 5 meters for 16 Gbps SFPs.

Buffers required

The number of buffers required for the given distance and speed. Irrespective of test frame size, the number of buffers required is calculated for 2112-byte frames.

Failure report

Displays details on the local port and remote port errors, if any, when the D_Port test fails.

Use the **--show all** command to display the following test summary information:

Port

The port number on the local side of the link

State

The port state: OFFLINE or ONLINE

SFP Capabilities

Electrical loopback (E) or optical loopback (O), or both (E, O)

Test Result

PASSED, FAILED, or STOPPED

Egress power

The power values and power loss in the transmit path. The following values are displayed: power transmitted from the local side of the link (**Tx**), power received on the remote side of the link (**Rx**), the power loss, that is, the relative difference in power (**Diff**), and the status of the power loss (No Loss, within tolerable limits, and outside tolerable limits).

Ingress power

The power values and power loss in the receive path. The following values are displayed: power received on the local side of the link (**Rx**), power transmitted from the remote side of the link (**Tx**), the power loss, that is, the relative difference in power (**Diff**), and the status of the power loss (No Loss, within tolerable limits, and outside tolerable limits).

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is supported only on Fibre Channel ports. SFPs must be Brocade-branded and run Fabric OS v7.0.0 or later. This command skips the optical loopback test for loopback plugs or when the port loops back to another D_Port in the same switch.

The port error statistics may show incremented values after D_Port tests are successful due to port toggle operation at the end of the D_Port test.

The approximate cable distance is calculated and displayed instead of estimated cable distance if the "electrical loopback" and "optical loopback" tests statuses are SKIPPED. The associated value of Approximate/Estimated cable distance will be UNKNOWN if the "electrical loopback" and "optical loopback" tests statuses are "SKIPPED" and the link distance is less than 150 meters.

The D_Port long duration test can be run on only one port at a time.

Operands

This command has the following operands:

slot

On bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port_list

Specifies one or more ports, relative to the slot on bladed systems. Use **switchShow** for a listing of valid ports. A port list should be enclosed in double quotation marks and can consist of the following:

- A single port, for example, "8" or "5/8" on blades systems.
- A port range where beginning and end port are separated by a dash, for example, "8-13" or "5/8-13" on blades systems. A port range cannot span multiple slots.
- A set of ports separated by commas, for example "3,5,7,8" or "5/3,5,7,8" on bladed systems.
- A wildcard "*" indicates all ports. The wildcard can be represented as "*", or "***".

--start

Initiates the diagnostics on the specified D_Ports.

-nframes *number*

Specifies the number of frames to send in millions. The default value is 1 for 1 million frames. The maximum value is 2147483647 million frames. If the number of frames exceeds the maximum value, the test proceeds with the initial 2147483647 million frames.

-time *time*

Specifies the time duration in HH:MM format for which the frame traffic test will run.

-framesize *size*

Specifies the size of test frames that are generated to run the test. The range for the framesize is 36 to 2112 bytes. The size of the frames can be specified in multiples of 4; otherwise the nearest higher multiple of 4 value will be taken as frame size. Default value is 1024.

-pattern *pat_name*

Specifies the name of the predefined pattern to be used in the payload. Valid patterns are BYTE_NOT, QUAD_NOT, WORD_NOT, BYTE_RAMP, QUAD_RAMP, WORD_RAMP, BYTE_LFSR, RANDOM, CRPAT, CSPAT, CHALF_SQ, CQTR_SQ, RDRAM_PAT, JCRPAT, JCJTPAT, JCSPAT, PRED_RAND, SMI_TEST, CJPAT, QUAD_NOTP, JSPAT, JTSPAT. Use the **dataTypeShow** command to get the list of frame patterns.

-payload *payload_pattern*

Specifies the user defined payload pattern in hexadecimal.

-fec

Enables the forward error correction (FEC) during D_Port test. This sub-option is not supported on D_Ports configured with Dense Wavelength Division Multiplexing (DWDM).

| | |
|-------------------------------------|--|
| -cr | Enables the credit recovery (CR) during D_Port test. |
| --setarg | Sets the D_Port test parameters and exits. The sub-options of --setarg are same as --start . This option is not supported for D_Port HBA connections. |
| --stop | Terminates the diagnostics on the specified D_Ports. Brocade recommends that you execute this command on both the local port and the remote port to properly terminate the D_Port tests. |
| --restart | Restarts the D_Port diagnostic tests with the configured parameters. This option is not supported for D_Port HBA connections. |
| -set [ewrap owrap] | Sets Electrical loopback (E_WRAP) or Optical loopback (O_WRAP) on Brocade-branded 16G/32G supported SFPs. |
| -clear [ewrap owrap all] | Checks if Electrical loopback (E_WRAP) or Optical loopback (O_WRAP) is set on the given SFP port and clears the setting. Specify all to clear both E_WRAP and O_WRAP settings. |
| --show | Displays runtime status and results of the diagnostics. |
| -detail | Displays the current detailed D_Port test results and statistics. This option provides extra details in case of test failure. |
| --exit | Exits the on-demand or Dynamic D_Ports to normal mode. The static D_Ports are retained in D_Port mode. |
| --help | Displays the command usage. |

Examples

To initiate the D_Port test on a single port:

```
switch:admin> portdporttest --start
```

To terminate the D_Port test on a single port:

```
switch:admin> portdporttest --stop
```

To set EWRAP on a port:

```
switch:admin> portdporttest --set ewrap 10/32
```

To set OWRAP on a port:

```
switch:admin> portdporttest --set owrap 10/32
```

To clear both EWRAP and OWRAP on a port:

```
switch:admin> portdporttest --clear all 10/32
```

To display the D_Port test results for a port while the test is in progress or after the test has completed:

```
switch:admin> portdporttest --show 16
D-Port Information:
=====
Port:                               16
Remote WWNN:                         10:00:00:27:f8:f0:21:60
Remote port index:                   8
Mode:                                 Manual
No. of test frames:                  1 Million
Test frame size:                     1024 Bytes
FEC (enabled/option/active): Yes/No/Yes
CR (enabled/option/active): Yes/No/No
Start time:                          Sun Sep 27 17:15:09 2015
End time:                             Sun Sep 27 17:15:27 2015
Status:                               PASSED
=====
Test          Start time  Result      EST(HH:MM:SS)  Comments
=====
Electrical loopback  17:15:11    PASSED      -----
Optical loopback    17:15:18    PASSED      -----
Link traffic test   17:15:23    PASSED      -----
=====
Roundtrip link latency:  1061 nano-seconds
Estimated cable distance: 1 meters
Buffers required:       1 (for 2112 byte frames at 32Gbps speed)
Egress power:          Tx: -1.0 dBm, Rx: -0.9 dBm, Diff: 0.0 dBm (No Loss)
Ingress power:         Rx: -1.3 dBm, Tx: -1.4 dBm, Diff: 0.0 dBm (No Loss)
```

To display D_Port summary information:

```
switch:admin> portdporttest --show all
Port  State      SFP Capabilities  Test Result
=====
24   ONLINE    E,O               PASSED
26   ONLINE    E,O               PASSED
33   OFFLINE   ---               FAILED
```

To display the D_Port test results for a port when the test has failed:

```
switch:admin> portdporttest --show 3/20
D_Port Information:
=====
Slot:                               3
Port:                                20
Remote WWNN:                         10:00:00:05:33:7e:69:c4
Remote port index:                   28
Mode:                                 Manual
No. of test frames:                  1 Million
Test frame size:                     1024 Bytes
FEC (enabled/option/active): Yes/No/No
```

```

CR (enabled/option/active): Yes/No/No
Start time:      Mon Apr  8 12:16:12 2013
End time:        Mon Apr  8 12:16:29 2013
Status:          FAILED
=====
Test            Start time  Result  EST(HH:MM:SS)  Comments
=====
Electrical loopback 12:16:13  PASSED  -----
Optical loopback   12:16:18  PASSED  -----
Link traffic test  12:16:29  FAILED  -----      Remote port is not ready to
start the test
=====
Roundtrip link latency: unknown

Failure report:
Remote port error(s): GENERIC, BAD_EOF, UNREACHABLE, SIGNAL_LOSS

Use "portdporttest --show -detail" for more details

```

To display the detailed D_Port test results for a port when the test has failed:

```

switch:admin> portdporttest --show -detail 1
*****
1. Current D-Port test results:
*****
Port: 1
Remote WNN:10:00:00:05:33:7e:97:e9
Remote port index:2
Mode: Automatic
No. of test frames:1 Million
Test frame size:1024 Bytes
FEC (enabled/option/active):Yes/No/No
CR (enabled/option/active):No/No/No
Start time:Wed Feb 12 05:57:41 2014
End time:Wed Feb 12 05:57:48 2014
Status:FAILED
=====
Test      Start timeResultEST(HH:MM:SS)Comments
=====
Electrical loopback05:57:42PASSED  -----
Optical loopback  -----FAILED  -----Remote port is not ready to start the
test
Link traffic test  -----NOT STARTED-----
=====
Roundtrip link latency:unknown

Failure report:

*****
2. Current Port statistics:
*****
portstatshow 1:
=====
stat_wtx      1000436935  4-byte words transmitted
stat_wrx      1000437023  4-byte words received
stat_ftx      3775241     Frames transmitted
stat_frx      3775241     Frames received
stat_c2_frx   0           Class 2 frames received
stat_c3_frx   3775233     Class 3 frames received
stat_lc_rx    4           Link control frames received

```

```

stat_mc_rx          0          Multicast frames received
stat_mc_to          0          Multicast timeouts
stat_mc_tx          0          Multicast frames transmitted
tim_rdy_pri         0          Time R_RDY high priority
tim_txcrd_z         3489       Time TX Credit Zero (2.5Us ticks)
tim_txcrd_z_vc 0- 3: 3489      0          0          0
tim_txcrd_z_vc 4- 7: 0        0          0          0
tim_txcrd_z_vc 8-11: 0        0          0          0
tim_txcrd_z_vc 12-15: 0       0          0          0
er_enc_in           0          Encoding errors inside of frames
er_crc              0          Frames with CRC errors
er_trunc            0          Frames shorter than minimum
er_toolong          0          Frames longer than maximum
er_bad_eof          0          Frames with bad end-of-frame
er_enc_out          0          Encoding error outside of frames
er_bad_os           0          Invalid ordered set
er_pcs_blk          0          PCS block errors
er_rx_c3_timeout    0          Class 3 receive frames discarded due to timeout
er_tx_c3_timeout    0          Class 3 transmit frames discarded due to timeout
er_unroutable       0          Frames that are unroutable
er_unreachable      0          Frames with unreachable destination
er_other_discard    0          Other discards
er_type1_miss       0          frames with FTB type 1 miss
er_type2_miss       0          frames with FTB type 2 miss
er_type6_miss       0          frames with FTB type 6 miss
er_zone_miss        0          frames with hard zoning miss
er_lun_zone_miss    0          frames with LUN zoning miss
er_crc_good_eof     0          Crc error with good eof
er_inv_arb          0          Invalid ARB
er_single_credit_loss0 0          Single vcrdy/frame loss on link
er_multi_credit_loss0 0          Multiple vcrdy/frame loss on link

```

porterrshow:

```

=====
          frames      enc   crc   crc   too   too   bad   enc   disc   link
loss  loss  frjt  fbsy  c3timeout  pcs
          tx    rx    in   err   g_eof  shrt  long  eof   out  c3   fail
sync  sig
1:   3.7m  3.7m  0    0    0    0    0    0    0    0    0
0    0    0    0    0    0

```

```

*****
3. Port statistics prior to the port was set to D-Port:
*****

```

portstatshow 1

```

=====
stat_wtx           353          4-byte words transmitted
stat_wrx           444          4-byte words received
stat_ftx           8          Frames transmitted
stat_frx           8          Frames received
stat_c2_frx        0          Class 2 frames received
stat_c3_frx        0          Class 3 frames received
stat_lc_rx         4          Link control frames received
stat_mc_rx         0          Multicast frames received
stat_mc_to         0          Multicast timeouts
stat_mc_tx         0          Multicast frames transmitted
tim_rdy_pri        0          Time R_RDY high priority
tim_txcrd_z        1892       Time TX Credit Zero (2.5Us ticks)
tim_txcrd_z_vc 0- 3: 1892      0          0          0
tim_txcrd_z_vc 4- 7: 0        0          0          0

```

```

tim_txcrd_z_vc 8-11: 0          0          0          0
tim_txcrd_z_vc 12-15: 0         0          0          0
er_enc_in      0          Encoding errors inside of frames
er_crc         0          Frames with CRC errors
er_trunc       0          Frames shorter than minimum
er_toolong     0          Frames longer than maximum
er_bad_eof     0          Frames with bad end-of-frame
er_enc_out     0          Encoding error outside of frames
er_bad_os      0          Invalid ordered set
er_pcs_blk     0          PCS block errors
er_rx_c3_timeout 0        Class 3 receive frames discarded due to timeout
er_tx_c3_timeout 0        Class 3 transmit frames discarded due to timeout
er_unroutable  0          Frames that are unroutable
er_unreachable 0          Frames with unreachable destination
er_other_discard 0        Other discards
er_type1_miss  0          frames with FTB type 1 miss
er_type2_miss  0          frames with FTB type 2 miss
er_type6_miss  0          frames with FTB type 6 miss
er_zone_miss   0          frames with hard zoning miss
er_lun_zone_miss 0        frames with LUN zoning miss
er_crc_good_eof 0        Crc error with good eof
er_inv_arb     0          Invalid ARB
er_single_credit_loss0 0        Single vcrdy/frame loss on link
er_multi_credit_loss0 0        Multiple vcrdy/frame loss on link

```

```
porterrshow
```

```
=====
```

| | frames | enc | crc | crc | too | too | bad | enc | disc | link | loss | |
|------|--------|------|-----------|-----|-------|------|------|-----|------|------|------|------|
| loss | frjt | fbsy | c3timeout | pcs | g_eof | shrt | long | eof | out | c3 | fail | sync |
| sig | tx | rx | in | err | g_eof | shrt | long | eof | out | c3 | fail | sync |
| | tx | rx | in | err | g_eof | shrt | long | eof | out | c3 | fail | sync |
| 1: | 4 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 11: | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 21: | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 31: | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 41: | 8.6k | 8.6k | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

SEE ALSO [dataTypeShow](#), [fabricLog](#), [portCfg](#), [portCfgDPort](#), [portShow](#), [switchShow](#)

portEnable

Enables a port or a range of ports.

Synopsis `portenable [slot/]port1[-port2] [...]`
`portenable -i [index1[-index2] [...]] [-f]`
`portenable -x [hex1[-hex2] [...]]`
`portenable -slot [slot1[-slot2]]`
`portenable -h`

Description Use this command to enable a port or a range of ports. If a port is connected to another switch when this command is issued, the fabric may reconfigure. After the port is enabled, devices connected to the port can again communicate with the fabric. The front panel LED of a port that is enabled and online is green.

You can identify a single port to be enabled by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers(decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

For ports that come online after being enabled, the following indications might be sent to indicate a state transition: RSCN, SNMP trap, Web pop-up window.

This command fails if the switch is disabled, the port's blade is not fully enabled (faulted, powered off, or disabled), or if the port is persistently disabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports. To enable an FCoE port, use **fcoe --enable**.

Operands This command has the following operands:

slot On bladed systems only, specifies the slot number of the ports to be enabled, followed by a slash (/).

port1[-port2] Enables a single port or a range of ports identified by port numbers. The port range cannot span slots, but you can specify multiple port ranges separated by a space, for example 3/1-4 4/7-9.

-i index1[-index2] Enables a port or a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, 33-38 40-60.

- f** Ignores nonexistent ports. This operand is valid only with the **-i** option.
- x [hex1 [-hex2]]** Enables a port or a range of ports identified by port numbers, index number in hexadecimal format. You may specify multiple port ranges separated by a space, for example, **-x 21-26 28-3c**.
- slot [slot1[-slot2]]** Enables all ports on a slot or on a range of slots, for example, **-s 3-5**. Multiple slot ranges are not supported with this command.
- h** Displays the command usage.

Examples

To enable a single port:

```
switch:admin> portenable 2/4
```

To enable a range of ports:

```
switch:admin> portenable 2/4-8
```

To enable multiple port ranges:

```
switch:admin> portenable 2/24-26 3/10-12 4/3-4
```

To enable a port by specifying its index number:

```
switch:admin> portenable -i 176
portenable: portSwap feature enabled.
```

To enable a range of ports by specifying the corresponding port index range:

```
switch:admin> portenable -i 170-176
portenable: portSwap feature enabled.
```

To enable multiple ports by specifying multiple port index ranges:

```
switch:admin> portenable -i 30-36 170-176
portenable: portSwap feature enabled.
```

To enable a range of ports by specifying port index number in hexadecimal format:

```
switch:admin> portenable -x 1d-1e
```

To enable all ports on slot 3-5.

```
switch:admin> portenable -s 3-5
```

SEE ALSO

[portCfgPersistentDisable](#), [portCfgPersistentEnable](#), [portDisable](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portEncCompShow

Displays encryption and compression port configuration details.

Synopsis `portenccompshow`

Description Use this command to display a list of ports that can be configured for encryption or compression. The command displays one section per ASIC and the ports configurable for that ASIC. The output includes the following information:

User Port

The port index number of the port to be configured. Use the **switchShow** command to identify the corresponding slot and port number.

Encryption configured

Displays "Yes" if encryption is enabled on the port. Displays "No" if encryption is disabled. This parameter is configured with the **portCfgEncrypt** command.

Encryption active

Displays the port's runtime status. "Yes" indicates that the port is online and enabled for encryption. "No" indicates that the port is enabled for encryption but offline, or not enabled for encryption.

Compression configured

Displays "Yes" if compression is enabled on the port. Displays "No" if compression is disabled. This parameter is configured with the **portCfgCompress** command.

Compression active

Displays the port's runtime status. "Yes" indicates that the port is online and enabled for compression. "No" indicates that the port is enabled for compression but offline, or not enabled for compression.

Speed

Displays the speed of the port and link speed of the online ports. If the speed is configured as AUTO NEG, the speed of the port is taken as the maximum speed of the port.

Notes On 16 Gbps-capable platforms, both encryption and compression features are supported. On 32 Gbps-capable platforms, only compression is supported and encryption is not supported.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the port configuration and runtime status of the ports configurable for encryption or compression on a DCX 8510-8:

```
switch:admin> portenccompshow
User      Encryption      Compression      Config
Port      Configured      Active           Configured      Active           Speed
-----
    17      No             No              No             No
```

| | | | | | |
|-------|-----|-----|-----|----|-----|
| 18 | No | No | No | No | |
| 19 | No | No | No | No | |
| 20 | No | No | No | No | |
| 21 | No | No | No | No | |
| 22 | No | No | No | No | |
| 23 | No | No | No | No | |
| 144 | Yes | Yes | Yes | No | 16G |
| 145 | No | No | No | No | |
| 146 | No | No | No | No | |
| 147 | No | No | No | No | |
| 148 | No | No | No | No | |
| 149 | Yes | No | Yes | No | 16G |
| 150 | No | No | No | No | |
| 151 | No | No | No | No | |
| ----- | | | | | |
| 88 | No | No | No | No | |
| 89 | No | No | No | No | |
| 90 | No | No | No | No | |
| 91 | No | No | No | No | |
| 92 | No | No | No | No | |
| 93 | No | No | No | No | |
| 94 | No | No | No | No | |
| 95 | No | No | No | No | |
| 208 | No | No | No | No | |
| 209 | No | No | No | No | |
| 210 | No | No | No | No | |
| 211 | No | No | No | No | |
| 212 | No | No | No | No | |

SEE ALSO [portCfgEncrypt](#), [portCfgCompress](#)

portErrShow

Displays a port error summary.

Synopsis

porterrshow

porterrshow *[[slot]port1[-[slot]port2]]*

porterrshow **-i** | **-index** *[index1[-index2]]*

porterrshow **-x** *[hex1[-hex2]]*

porterrshow **--help**

Description

Use this command to display an error summary for a port or a range of ports. Counts are reported on frames transmitted by the port (Tx) or on frames received by the port (Rx). The display contains one output line per port. Numeric values exceeding 999 are displayed in units of thousands (k), millions (m), or giga (g) if indicated.

You can identify a single port to be configured by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers (decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers. When used without operands, this command displays error summary for all ports on the switch.

Specifying multiple ports with the index number (**-i** or **-x**) or slot (**-s**) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the **-i** option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

Values for the following parameters are displayed:

| | |
|------------------|--|
| frames tx | Number of frames transmitted (Tx). |
| frames rx | Number of frames received (Rx). |
| enc in | Number of encoding errors inside frames received (Rx). |
| crc err | Number of frames with CRC errors received (Rx). |
| crc g_eof | Number of frames with CRC errors with good EOF received (Rx). |
| too shrt | Number of frames shorter than minimum received (Rx). |
| too long | Number of frames longer than maximum received (Rx). |
| bad eof | Number of frames with bad end-of-frame delimiters received (Rx). |
| enc out | Number of encoding error outside of frames received (Rx). |

| | |
|----------------------|---|
| disc c3 | Number of Class 3 frames discarded (Rx). This counter includes the sum of the following class 3 discard counters reported by the portStatsShow command: er_rx_c3_timeout , er_tx_c2_timeout , er_c2_dest_unreach , and er_other_disc . Refer to portStatsShow help for a description of these counters. |
| link fail | Number of link failures (LF1 or LF2 states) received (Rx). |
| loss sync | Number of times synchronization was lost (Rx). |
| loss sig | Number of times a loss of signal was received (increments whenever an SFP is removed) (Rx). |
| frjt | Number of transmitted frames rejected with F_RJT (Tx). |
| fbsy | Number of transmitted frames busied with F_BSY (Tx). |
| c3-timeout tx | The number of transmit class 3 frames discarded at the transmission port due to timeout (platform- and port-specific). |
| c3-timeout rx | The number of receive class 3 frames received at this port and discarded at the transmission port due to timeout (platform- and port-specific). |
| pcs err | The number of Physical Coding Sublayer (PCS) block errors. This counter records encoding violations on 10 Gbps or 16 Gbps ports. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports.

Operands This command has the following optional operands:

[slot]port1[-[slot]port2]

Displays port error summary for a single port or for a range of ports, relative to the slot number on bladed systems. Port numbers in a range must be separated by a dash (-), for example, 3-5, or 2/0-2/15. Port ranges cannot span slots. Use **switchShow** to display a listing of valid ports. Port operands are optional; if omitted, information for all ports is displayed.

-i | -index [index1[-index2]]

Displays port error summary for a single port or for a range of ports identified by port index numbers. You may specify index ranges separated by "-", for example, -i 33-47.

portErrShow

-x [*hex1* [*-hex2*]]

Displays port error summary for a port or a range of ports identified by port index numbers in hexadecimal format. You may specify port ranges separated by "-", for example, -x 21-26.

--help

Displays the command usage.

Examples

To display error counters for ports on a switch:

```
switch:admin> porterrshow
      frames enc   crc   crc   too   too   bad   enc \
      tx    rx    in   err   g_eof shrt  long eof  out \
===== \
0:    0      0      0     0     0     0     0     0     0 \
1:  2.5g  73m      0     0     0     0     0     0     1.2k \
2:    0      0      0     0     0     0     0     0     0 \
3:    0      0      0     0     0     0     0     0     667 \

      disc  link  loss  loss  frjt  fbsy  c3-timeout  pcs
      c3   fail sync  sig           tx    rx    err
=====
      0      0      2      4      0      0      0      0      0
      563.2k 0      0      0      1      0      0    415.2k  0
      0      0      0      1      0      0      0      0      0
      0     21      4      8      0      0      0      0      0
(output truncated)
```

SEE ALSO [portShow](#), [portStatsShow](#)

portFlagsShow

Displays the port status bitmaps for all ports in a switch.

Synopsis `portflagsshow`

Description Use this command to display the following status for a port:

SNMP

Displays whether the port is online or offline.

Physical

Displays the port physical status. Valid values are In_Sync, No_Light (8 Gbps-capable ports only), and No_SigDet (quad small form-factor pluggables (QSFPs) installed without cables). Refer to **switchShow** help for more information on these states.

Flags

Displays whether there is an SFP inserted in the port, whether the port is active, and the port type.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the port status for all ports in the switch:

```
switch:user> portflagsshow
Port SNMP      Physical      Flags
-----
 0 Online      In_Sync      PRESENT ACTIVE F_PORT G_PORT \
U_PORT LOGICAL_ONLINE LOGIN NOELP ACCEPT FLOGI
 1 Online      In_Sync      PRESENT ACTIVE E_PORT G_PORT \
U_PORT SEGMENTED CBL_LB LOGIN LED
 2 Offline     No_Module    PRESENT U_PORT LED
 3 Offline     No_Light     PRESENT U_PORT LED
 4 Offline     No_Module    PRESENT U_PORT LED
 5 Offline     No_Module    PRESENT U_PORT LED
 6 Offline     No_Module    PRESENT U_PORT LED
 7 Offline     No_Module    PRESENT U_PORT LED
 8 Offline     No_Module    PRESENT U_PORT LED
 9 Offline     No_Module    PRESENT U_PORT LED
10 Online      In_Sync      PRESENT ACTIVE F_PORT G_PORT \
U_PORT LOGICAL_ONLINE LOGIN NOELP ACCEPT FLOGI
11 Offline     No_Module    PRESENT U_PORT LED
12 Offline     No_Module    PRESENT U_PORT LED
13 Offline     No_Module    PRESENT U_PORT LED
14 Offline     No_Module    PRESENT U_PORT LED
15 Offline     No_Module    PRESENT U_PORT LED
16 Online      In_Sync      PRESENT ACTIVE F_PORT G_PORT \
U_PORT LOGICAL_ONLINE LOGIN NOELP ACCEPT FLOGI
17 Online      In_Sync      PRESENT ACTIVE F_PORT G_PORT \
U_PORT LOGICAL_ONLINE LOGIN NOELP ACCEPT FLOGI
18 Offline     No_Module    PRESENT U_PORT LED
```

portFlagsShow

```
19 Offline   No_Module  PRESENT U_PORT LED
20 Offline   No_Module  PRESENT U_PORT LED
21 Offline   No_Module  PRESENT U_PORT LED
22 Offline   No_Module  PRESENT U_PORT LED
(output truncated)
```

To display the port status for QSFPs installed without cables (partial output) :

```
6  16 Offline   No_SigDet  PRESENT U_PORT LED
6  17 Offline   No_SigDet  PRESENT U_PORT LED
6  18 Offline   No_SigDet  PRESENT U_PORT LED
6  19 Offline   No_SigDet  PRESENT U_PORT LED
```

To display the port status for QSFPs installed and connected with cables (partial output) :

```
6  28 Online     In_Sync    PRESENT ACTIVE E_PORT \
T_PORT T_MASTER G_PORT U_PORT LOGICAL_ONLINE LOGIN LED
6  29 Online     In_Sync    PRESENT ACTIVE E_PORT T_PORT \
T_MASTER G_PORT U_PORT LOGICAL_ONLINE LOGIN LED
6  30 Online     In_Sync    PRESENT ACTIVE E_PORT \
T_PORT T_MASTER G_PORT U_PORT
```

SEE ALSO [portShow](#), [switchShow](#)

portLedTest

Cycles user port LEDs.

Synopsis `portledtest`
 `[--slot slot_number]`
 `[-ports itemlist]`
 `[-uports itemlist]`
 `[-npass count]`
 `[-action value]`

Description Use this command to exercise the user port LEDs in the current switch. When used without a port specifier, all user ports are tested.

For each port, the command cycles through the ON and OFF state by setting the ATTN LEDs to green for the ON condition and unlighted for the OFF condition. The SPEED LEDs are initially set to black before the test starts. The SPEED LEDs turn green while the test is running.

This diagnostic cannot be run on an operational switch. You must disable the switch using the **chassisDisable** command before you can run this test. After the command completes, the ATTN LEDs flash amber, indicating that the command has finished and exited. Enable the switch using the **chassisEnable** command to set the ATTN LEDs back to black.

The port LED behavior varies depending on the platform.

Notes You cannot interrupt the test by pressing the return key (<cr>).

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--slot *slot* Specifies the slot number on bladed systems. When used without the **port** option, this option exercises all user ports in the specified slot. When used with the **port** option, only the ports specified for the specified slot are tested.

-ports *itemlist* Specifies a list of front-end blade ports to test. When a specific port is specified, the blade port reference is used. If you want to test a specific user port, you will need root access to the **bladePortMap** command to map the specified blade port to the desired user port. You can also use the **-uports** option to specify the user ports. Refer to **itemList** for more information on the **itemlist** parameter. When specifying ports on a chassis, the **-ports** operand must be preceded by the **--slot** operand.

-uports *itemlist* Specifies a list of user ports to test. If this operand is omitted, by default the test will run on all valid blade ports in the slot. Refer to **itemList** for more information on the **itemlist** parameter. This option is only available for 16 Gbps-capable blades.

-npass *count* Specify the number of times to perform this test. The default value is 10.

| | |
|------------------------------|--|
| -action <i>action</i> | Specifies the LED color. Valid values include the following: |
| 0 | This is the default action. Cycle all Port LEDs. C3 cycles are faster than C2 cycles. |
| 1 | Turn Port status LED off. |
| 2 | Turn Port status LED amber. |
| 3 | Turn Port status LED green. |
| 4 | Blink Port status LED Green. Supported only on C4 platforms. |
| 5 | Blink Port status LED Amber. Supported only on C4 platforms. |
| 16 | Not supported for C2, C3, and C4 platforms. |
| 17 | Not supported for C2, C3, and C4 platforms. Supported for the C3 core, and the port status LED blinks amber. |

Examples To test port LEDs on user port 1 use the corresponding value in the "BPt" column of the **bladePortMap** command (requires root permission):

```
switch:admin> portledtest --slot 1 -ports 213

Running portledtest .....
WARNING:
This test should NOT be aborted in the middle. If aborted, current
blade or the switch (in-case of Pizza box) may become unusable.
Reset the blade or the switch to recover.

PASSED.
```

SEE ALSO [chassisDisable](#), [chassisEnable](#), [fcipLedTest](#), [itemList](#), [switchDisable](#), [switchEnable](#)

portLogClear

Clears the port log.

Synopsis `portlogclear`

Description Use this command to clear the port log. It is recommended that you clear the port log before triggering an activity so that the log displays only the log events related to that activity.

If the port log is disabled, **portLogClear** enables it. The port log is disabled automatically when certain errors occur to allow the collection of all the information needed to understand the cause of the error. When the port log is disabled, the events already present in the log are preserved, but new events are not collected.

The following errors disable the port log:

- FCPH, EXCHBAD
- FCPH, EXCHFEE
- NBFMS, DUPEPORTSCN
- UCAST, RELICPDB

Refer to the *Fabric OS Message Reference* for more information on these errors.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To clear the port log:

```
switch:admin> portlogclear
switch:admin> portlogshow
port log is empty
```

SEE ALSO [portLogDump](#), [portLogShow](#)

portLogConfigShow

portLogConfigShow

Displays the current port log configuration.

Synopsis `portlogconfigshow`

Description Use this command to display the current port log configuration.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the current port log configuration:

```
switch:admin> portlogconfigshow
max portlog entries = 16384
```

SEE ALSO [portLogResize](#)

portLogDisable

Disables the port log facility.

Synopsis `portlogdisable`

Description Use this command to disable the port log facility.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the port log facility:

```
switch:admin> portlogdisable
```

SEE ALSO [portLogEnable](#)

portLogDump

Displays the port log without page breaks.

Synopsis `portlogdump [count[, saved]]`

Description Use this command to display the port log, listing all entries in the log without page breaks. This command displays the same information as **portLogShow**, but **portLogShow** prompts you to press Enter between each page.

For an explanation of the information displayed by this command, refer to the **portLogShow** command.

If the port log is disabled while this command is executed, a warning message is displayed. Refer to the **portLogClear** command for more information.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

count

Specifies the maximum number of lines to be displayed. Only the most recent count entries are displayed. This operand is optional.

saved

Specify a nonzero value to display the saved port log from the last switch fault. Refer to **upTime** for conditions that cause a fault. The operand count is ignored when displaying the saved log. This operand is optional.

Examples To display 10lines of the portlog on a Brocade DCX:

```
switch:user> portlogdump 10
time          task          event    port cmd  args
-----
Mon Nov 16 21:52:15 2009
21:52:15.214 FCPH          seq      106  7f0
ed210000,00000000,000073ee,10010082,00008000
21:52:15.214 PORT          Tx3      106  2032 \
02ffffffd,00ffffffd,0701ffff,13010000
21:52:15.216 FCPH          read     106  2032
03ffffffd,00ffffffd,00000000,00008000,07010000
21:52:15.216 FCPH          seq      106  7f0
ed980000,07010000,00004143,0004001c,00008000
21:52:15.216 FCPH          write    106  2032
00ffffffd,00ffffffd,00000000,00008000,00000000
21:52:15.216 FCPH          seq      106  7f0
ed210000,00000000,000073ee,10010082,00008000
21:52:15.216 PORT          Tx3      106  2032 \
02ffffffd,00ffffffd,06ffffff,13010000
21:52:15.218 FCPH          read     106  2032
```

```
03ffffffd,00ffffffd,00000000,00008000,06fe0000
21:52:15.218 FCPH      seq      106  7f0 \
ed980000,06fe0000,00004143,0004001c,00008000
21:52:15.220 FCPH      seq      106  7f0
ed980000,06fc0000,00004143,0004001c,00008000
```

SEE ALSO [portLogClear](#), [portLogShow](#), [upTime](#)

portLogDumpPort

Displays the port log of a specified port without page breaks.

Synopsis `portlogdumpport port_index`

Description Use this command to display the port log for a single port specified by its port index number. The command displays all entries in the log without any page breaks. This command is identical to **portLogShowPort**, except that **portLogShowPort** prompts you to press Enter to display the next page.

Port logs are circular log files in the switch firmware, which can save up to 65,536 entries depending on the hardware platform. Use **portLogConfigShow** to display the current size of the port log. Once the log has reached its maximum size, new entries displace the oldest ones. Port logs capture switch-to-device, device-to-switch, switch-to-switch, some device-to-device1, and control information

If the port log is disabled while this command is executed, a warning message is displayed. Refer to **portLogClear** command for more information.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

port_index

Displays the port log for the a single port specified by its port index number. Use **switchShow** for a listing of valid port index numbers.

Examples To display the port log dump for a port:

```
switch:user> portlogdumpport 14
time          task          event port  cmd  args
-----
08:35:27.899  tShell           pstate 14   OL1
08:35:27.899  tReceive         pstate 14   LR2
08:35:27.916  tReceive         pstate 14   AC
08:35:28.416  interrupt        scn    14   1
08:35:28.433  tFabric          ioctl  14   90  \
101d9910,0
08:35:28.433  tFabric          Tx     14  164  \
02ffffffd,00ffffffd,0005ffff,10000000
08:35:28.433  tReceive         Rx     14   0  \
c0ffffffd,00ffffffd,00050006
08:35:28.433  tReceive         Rx     14  164  \
03ffffffd,00ffffffd,00050006,02000000
08:35:28.433  tTransmit        Tx     14   0  \
c0ffffffd,00ffffffd,00050006
08:35:28.433  tFabric          ioctl  14   91  \
103646d8,0
08:35:28.466  tFabric          ioctl  14   a7  3c,1
(output truncated)
```

SEE ALSO [portLogDump](#), [portLogClear](#), [portLogShow](#), [upTime](#)

portLogEnable

Enables the port log facility.

Synopsis `portogenable`

Description Use this command to enable the port log facility.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable the port log facility:

```
switch:admin> portlogenable
```

SEE ALSO [portLogDisable](#)

portLogEventShow

Displays information about port log events.

Synopsis `portlogeventshow`

Description Use this command to display information about the ID associated with the various port log events. The Disabled field indicates whether the port log for that event ID is disabled (1) or enabled (0).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display information about port log events:

```
switch:admin> portlogeventshow
ID      Event-Name      Disabled
-----
1       start           0
2       disable         0
3       enable          0
4       ioctl           0
5       Tx              0
6       Tx1             0
7       Tx2             0
8       Tx3             0
9       Rx              0
10      Rx1             0
11      Rx2             0
12      Rx3             0
13      stats           0
14      scn             0
15      pstate         0
16      reject         0
17      busy           0
18      ctin           0
19      ctout          0
20      errlog         0
21      loopscn        0
22      create         0
23      debug          1
24      nbrfsm         0
25      timer          0
(output truncated)
```

SEE ALSO [portLogTypeDisable](#), [portLogTypeEnable](#)

portLoginShow

Displays port login status of devices attached to the specified port and the details of the last logged out device.

Synopsis `portloginshow [slot/port] [-history]`

Description Use this command to display port login status received from devices attached to the specified port. For each login, this command displays the following fields:

| | |
|----------------------|---|
| Type | Type of login can display one of the following: |
| fd | FDISC, Discover F_Port Service Parameters or Virtual N_Port login. |
| fe | FLOGI, Fabric Login to Fabric F_Port. |
| ff | PLOGI, Port Login to specific N_Ports or well-known addresses like Name Server. |
| PID | The 24-bit Port ID of the attached device. |
| WorldWideName | The port's world wide name. |
| credit | The credit for this login as appropriate. This is BB (buffer-to-buffer) credit for Flogs and EE (end-to-end) credit for PLOGIs. |
| df_sz | The default frame size for this login. |
| cos | Class of Services supported. This can be a combination of the following bits: |
| 4 | Class 2 is supported. |
| 8 | Class 3 is supported. |
| logout time | Timestamp when the device is logged out from a port. |

Further information about each login is displayed after these columns, including the Port ID of the well-known address or N_Port that was the target of the PLOGI, if applicable.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not applicable to embedded FCoE ports. Use **fcoeLoginShow**.

portLoginShow

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the port to be displayed, followed by a slash (/).

port

Specify the port for which to display login status information, relative to its slot for bladed systems. Use **switchShow** for a list of valid ports.

-history

Displays the details of the last logged out device on a port for the login types "fe" and "fd". The time stamp displayed under the column "logout time" is in the UTC format.

Examples To display the logins received by Port 23 (revealing one FLOGI (type fe) and two PLOGIs):

```
switch:admin> portloginshow 23
Type  PID      World Wide Name      credit df_sz cos
=====
fe  201700  21:00:00:e0:8b:05:a3:c9  3  2048   8  scr=1
ff  201700  21:00:00:e0:8b:05:a3:c9  0    0    8  d_id=FFFC20
ff  201700  21:00:00:e0:8b:05:a3:c9  0    0    8  d_id=FFFFFC
```

To display the last logged out device on a port:

```
switch:admin> portloginshow 1/0 -history
Type  PID      World Wide Name      logout time
=====
fd  550002  30:0c:02:05:1e:61:23:8f  09/17/2014 13:56:01
fd  550001  30:0c:01:05:1e:61:23:8f  09/17/2014 13:56:02
fe  550000  30:0c:00:05:1e:61:23:8f  09/17/2014 13:56:02
fd  550001  30:0c:01:05:1e:61:23:8f  09/18/2014 05:49:37
```

SEE ALSO [fcpProbeShow](#), [portShow](#)

portLogPdisc

Sets or clears the debug_pdisc_flag.

Synopsis `portlogpdisc 0 | 1`

Description Use this command to set or clear the debug_pdisc_flag. This command is part of the environmental monitor. A setting of 1 enables logging of Port Discovery parameters. The PDISC log is disabled by default.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

`0 | 1`
Specify 0 to clear or 1 to set the debug_pdisc_flag. The default is 0.

Examples To set the debug_pdisc_flag:

```
switch:admin> portlogpdisc 1
PDISC log setting = 1
```

SEE ALSO `None`

portLogReset

portLogReset

Enables the port log facility.

Synopsis `portlogreset`

Description Use this command to enable the port log facility.

Notes Refer to **portLogClear** for events that may disable the port log facility.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable the port log:

```
switch:admin> portlogreset
```

SEE ALSO None

portLogResize

Resizes the port log to include a specified number of entries.

Synopsis `portlogresize num_entries`

Description Use this command to resize the port log to include a specified number of entries. If the specified number of entries is less than the already configured port log size, there is no change.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

num_entries

Specifies the number of port log entries. The valid range of values is 16,384 to 32,768.

Examples To resize the portlog:

```
switch:admin> portlogresize 17288
```

SEE ALSO [portLogConfigShow](#)

portLogShow

Displays the port log with page breaks.

Synopsis `portlogshow [count[, saved]]`

Description Use this command to display the port log with page breaks. This command displays the same information as **portLogDump**, but one page at a time.

The port log is a circular log file in the switch firmware which can save up to 65,536 entries depending on the hardware platform. Use **portLogConfigShow** to display the current port log size. Once the log has reached the maximum size, new entries replace the oldest ones. The port log captures switch-to-device, device-to-switch, switch-to-switch, some device-to-device, and control information.

If the command is executed while the port log is disabled, a warning message is displayed. Refer to the **portLogClear** command for more information.

For each log entry, the following information is displayed:

| | |
|----------------|---|
| Time | Displays the event date and time in milliseconds. The clock resolution is 16 milliseconds. |
| Task | Displays the name of the task that logged the event or "interrupt" if the event was logged in interrupt context, or "unknown" if the task no longer exists. |
| Event | Displays the task event that generated the log entry. Possible events include the following: |
| start | A switch start or restart event. |
| disable | A port is disabled. |
| enable | A port is enabled. |
| ioctl | A port I/O control is executed. |
| Tx | A frame is transmitted (class is indicated). |
| Rx | A frame is received (class is indicated). |
| scn | A state change notification is posted. |
| pstate | A port changes physical state. |
| reject | A received frame is rejected. |

| | |
|----------------|--|
| busy | A received frame is busy. |
| ctin | A CT based request is received. |
| ctout | A CT based response is transmitted. |
| errlog | A message is added to the error log. |
| loopscn | A loop state change notification is posted. |
| create | A task is created. |
| debug | Indicates a debug message. |
| nbrfsm | Indicates a neighbor state transition. |
| sn | Indicates a speed negotiation state. |
| fcin | Indicates an incoming Fibre Channel information unit. |
| fcout | Indicates an outgoing Fibre Channel information unit. |
| read | Indicates an information unit header log from a read operation. |
| write | Indicates an information unit header log from a write operation. |
| err | Indicates an information unit header log of an FC error frame. |
| frame | Indicates a Fibre Channel frame payload. |
| nsRemQ | Indicates an interswitch name server query. |
| rscn | Indicates a Registered State Change Notification. |
| xalloc | Allocates an exchange. |
| xfree | Frees an exchange. |
| xerr | Indicates an exchange error. |

portLogShow

| | |
|---------|---|
| xstate | Indicates an exchange state. |
| payload | Indicates a frame payload. |
| Port | Displays the port number that logged the event. |
| Cmd | Defined by the event. Displays a value defined by the event as follows: |
| ioctl | I/O control command code. |
| Tx & Rx | Frame payload size. |
| scn | New state (see state codes below). |
| pstate | New physical state (see pstate codes below). |
| ctin | The CT-subtype: |
| fc | Simple Name Server. |
| f8 | Alias Server. |
| ctout | The same as ctin. |
| errlog | Error level (refer to errShow). |
| loopscn | The current loop state during loop initialization. Possible values areas follows: |
| OLP | Offline (disconnected or nonparticipating). |
| LIP | FL_Port entered INITIALIZING or OPEN_INIT state. |
| LIM | LISM completed, FL_Port became the loop master. |
| BMP | Loop init completed, FL_Port in MONITORING state. |
| OLD | Port transitioned to the OLD_PORT state. |
| TMO | Loop init times out. |

| | |
|--------------------|--|
| Args | Displays additional information about the event as follows: |
| start | Start type: 0 = enable ports, 100 = disable ports. |
| disable | State (refer to state codes). |
| enable | Mode: 0 normal; nonzero loopback. |
| Tx & Rx | Header words 0,1,4 (R_CTL,D_ID,S_ID,OX_ID,RX_ID) and the first payload word. |
| reject | FC-PH reject reason. |
| busy | FC-PH busy reason. |
| ctin | <p>Argument 0 is divided into two 16-bit fields: [A] A bit map indicating whether subsequent arguments are valid (0001 means argument 1 is valid, 0003 means arguments 1 and 2 are valid). [B] The CT-based service command code.</p> <p>Argument 1 is the first word of the CT payload, if applicable (as specified in [A]). Argument 2 is the second word of the CT payload, if applicable (as specified in [A]).</p> |
| ctout | <p>Argument 0 is also divided into two 16-bit fields: [A] a bit map indicating whether subsequent arguments are valid (0001 means argument 1 is valid, 0003 means arguments 1 and 2 are valid). [B] the CT command code indicating whether an accept (8002) or a reject (8001). If [B] is an accept, argument 1 and 2 represents the first and second words of the CT payload, if applicable (as specified in [A]). If [B] is a reject, argument 1 contains the CT reject reason and explanation code.</p> |
| errlog | Error type (refer to errShow). |
| loopscn | The meaning further depends on each loop state: |
| OLP | Offline reason code, usually zero. |
| LIP | Reason code for LIPs initiated by FL_Port, if the code value is 800x (x = [1,0xc], see below), or the lower two bytes of the LIP received, if the code value is other than 800x. |
| LIM | Usually zero. |
| BMP | Memory address for the loop bitmap. |

| | |
|-----|--|
| OLD | Usually zero. |
| TMO | Encoded value of the state when loop initialization timed out This value is usually equal to the first word of a loop init frame payload. Other possible values include the following: |
| 2 | LIP (req. INITIALIZING) timeout. |
| 94 | FOFO ARB(FO) timeout. |
| 40 | CLS timeout. |

Codes used in various fields are as follows:

state Valid state values include the following:

| | |
|---|-----------|
| 1 | Online |
| 2 | Offline |
| 3 | Testing |
| 4 | Faulty |
| 5 | E_Port |
| 6 | F_Port |
| 7 | Segmented |

pstate Valid **pstate** values include the following:

| | |
|-----|----------------------------------|
| AC | Active State |
| LR1 | Link Reset: LR Transmit State |
| LR2 | Link Reset: LR Receive State |
| LR3 | Link Reset: LRR Receive State |
| LF1 | Link Failure: NOS Transmit State |

LF2 Link Failure: NOS Receive State

OL1 Offline: OLS Transmit State

OL2 Offline: OLS Receive State

OL3 Offline: Wait for OLS State

LIP reason

Valid **LIP reason** codes include the following:

8001 Retry loop init.

8002 Start loop after gaining sync.

8003 Restart loop after port reset.

8004 LIP when a loop hangs.

8005 Restart loop if LIP received when sending out ARB(FO).

8006 LIP when an OPN returns.

8007 Restart loop when LIPs received in OLD_PORT AC state.

8008 Restart loop if loop not empty but E_Port loopback.

8009 LIP as requested by the LINIT ELS received.

800a LIP as requested by the LPC ELS received.

Speed Negotiation States

Valid states include the following:

INIT Start negotiation.

NM Negotiate master.

WS Wait for signal.

NF Negotiation follows.

NC

Negotiation complete.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

count

Specifies the maximum number of lines to display. Only the most recent count entries are displayed. This operand is optional.

saved

Specifies a nonzero value to display the saved port log from the last switch fault. Refer to **upTime** for a list of conditions that cause a fault. The count is ignored when the saved log is displayed. This operand is optional.

Examples To view the port log for a port:

```
switch:user> portlogshow 24
time          task event port cmd  args
-----
17:05:30.384 PORT Rx      0  40  02ffffffd,00ffffffd,08fbffff,14000000
17:05:30.384 PORT Tx      0   0  c0ffffffd,00ffffffd,08fb0e02
17:05:30.384 PORT debug   0    0  00c0fffee,00fd0118,00000000,00000001
17:05:30.389 PORT Rx      1  40  02ffffffd,00ffffffd,08fdffff,14000000
17:05:30.389 PORT Tx      1   0  c0ffffffd,00ffffffd,08fd0e03
17:05:30.389 PORT debug   1    0  00c0fffee,00fd013c,00000000,00000001
17:05:30.504 PORT Rx      2  40  02ffffffd,00ffffffd,08feffff,14000000
17:05:30.504 PORT Tx      2   0  c0ffffffd,00ffffffd,08fe0e04
17:05:30.504 PORT debug   2    0  00c0fffee,00fd0182,00000000,00000001
17:05:30.507 PORT Rx      3  40  02ffffffd,00ffffffd,08ffffff,14000000
17:05:30.507 PORT Tx      3   0  c0ffffffd,00ffffffd,08ff0e05
17:05:30.508 PORT debug   3    0  00c0fffee,00fd0148,00000000,00000001
17:05:31.081 PORT Tx      0  40  02ffffffd,00ffffffd,0e06ffff,14000000
17:05:31.082 PORT debug   0    0  00c0fffee,00fd0188,14000000,00000001
17:05:31.084 PORT Rx      0   0  c0ffffffd,00ffffffd,0e060902
17:05:31.772 PORT Tx      1  40  02ffffffd,00ffffffd,0e07ffff,14000000
17:05:31.772 PORT debug   1    0  00c0fffee,00fd014a,14000000,00000001
17:05:31.774 PORT Rx      1   0  c0ffffffd,00ffffffd,0e070906
17:05:31.775 PORT Tx      2  40  02ffffffd,00ffffffd,0e08ffff,14000000
17:05:31.775 PORT debug   2    0  00c0fffee,00fd015c,14000000,00000001
17:05:31.777 PORT Rx      2   0  c0ffffffd,00ffffffd,0e080907
17:05:31.778 PORT Tx      3  40  02ffffffd,00ffffffd,0e09ffff,14000000
17:05:31.779 PORT debug   3    0  00c0fffee,00fd015e,14000000,00000001
17:05:31.782 PORT Rx      3   0  c0ffffffd,00ffffffd,0e090908
```

SEE ALSO [portLogClear](#), [portLogDump](#), [upTime](#)

portLogShowPort

Displays the port log of a specified port with page breaks.

Synopsis `portlogshowport port_index`

Description Use this command to display the port log of a specified port with page breaks. This command displays the same information as **portLogDumpPort**, except that **portLogDumpPort** does not prompt you to press Enter to display the next page.

If the command is executed while the port log is disabled, a warning message is displayed. Refer to the **portLogClear** command for more information.

Notes Refer to the **portLogShow** command for a description of the data returned by this command.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

port_index

Displays the port log for a single port specified by its port index number. Use **switchShow** for a listing of valid port index numbers.

Examples To display the port log for port 14:

```
switch:user> portlogshowport 14
time          task          event port  cmd  args
-----
08:35:28.483  tFabric        scn      14    0
08:35:27.899  tShell         pstate  14    OL1
08:35:27.899  tReceive       pstate  14    LR2
08:35:27.916  tReceive       pstate  14    AC
08:35:28.416  interrupt      scn      14    1
08:35:28.433  tFabric        ioctl    14    90  101d9910,0
08:35:28.433  tFabric        Tx       14    164  \
          02fffffd,00fffffd,0005ffff,10000000
08:35:28.433  tReceive       Rx       14    0  \
          c0fffffd,00fffffd,00050006
08:35:28.433  tReceive       Rx       14    164  \
          03fffffd,00fffffd,00050006,02000000
08:35:28.433  tTransmit      Tx       14    0  \
          c0fffffd,00fffffd,00050006
08:35:28.433  tFabric        ioctl    14    91  103646d8,0
08:35:28.433  tFabric        ioctl    14    92  103646d8,0
08:35:28.466  tFabric        ioctl    14    a7  3c,1
08:35:28.466  tFabric        pstate  14    LR1
08:35:28.466  tReceive       pstate  14    LR3
08:35:28.466  tReceive       pstate  14    AC
(output truncated)
```

SEE ALSO [portLogClear](#), [portLogDumpPort](#), [portLogShow](#), [upTime](#)

portLogTypeDisable

Disables the port log of a specified type.

Synopsis `portlogtypedisable id`

Description Use this command to disable the port log for a specified port log type.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

id

Specifies a nonzero value that corresponds to the port log type to be disabled. Use `portLogEventShow` for a listing of values corresponding to supported log types.

Examples To disable logging of type 2 port log events:

```
switch:admin> portlogtypedisable 2
```

SEE ALSO [portLogDisable](#), [portLogEventShow](#), [portLogTypeEnable](#)

portLogTypeEnable

Enables the port log of a specified port log type.

Synopsis `portlogtypeenable id`

Description Use this command to enable the port log for a specified port log type.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

id

Specifies a nonzero value that corresponds to the port log type to be enabled. Use `portLogEventShow` for a listing of values corresponding to supported log types.

Examples To enable logging of type 2 port log events:

```
switch:admin> portlogtypeenable 2
```

SEE ALSO [portLogEventShow](#), [portLogTypeDisable](#)

portLoopbackTest

Performs a functional test of port N->N path.

Synopsis `portloopbacktest`
 `[--slot slot]`
 `[-nframes count]`
 `[-lb_mode mode]`
 `[-spd_mode mode]`
 `[-ports itemlist]`
 `[-uports itemlist]`

Description Use this command to verify the functional operation of the switch by exercising the blade ports of the switch.

This test sends frames from a given port's transmitter and loops them back into the same port's receiver. The loopback is done at the parallel loopback path. The path traversed in this test does not include the media or the fiber cable. Only one frame is transmitted and received at any given time.

The port LED blinking pattern depends on the hardware platform on which the test is run. On 8 Gbps-capable platforms, the LEDs flicker green rapidly while the test is running. Only the ports on which the test is running flicker green. On 16 Gbps-capable platforms, the LEDs on all ports will flash amber, then turn to green, and then return to amber. After the test completes the lights turn to the expected switch disabled pattern. This difference in LED patterns is due to a difference in the ASIC design. Refer to the "Supported hardware and software" section in the *Fabric OS Command Reference* for a mapping of ASIC types to switch models.

The test performs the following operations:

1. Sets all ports for parallel loopback.
2. Creates a frame F of maximum data size (2,112 bytes).
3. Transmits the frame F through port N.
4. Picks up the frame from the same port N.
5. Checks if any of the following eight statistic error counters report nonzero values:
 ENC_in, CRC_err, TruncFrm, FrmTooLong, BadEOF, Enc_out, BadOrdSet, DiscC3
6. Checks whether the transmit, receive, or class 3 receiver counters are stuck at some value.
7. Checks whether the number of frames transmitted is not equal to the number of frames received.
8. Repeats Steps two through seven for all ports until one of the following conditions is met:
 - a. The number of frames (or pass count) requested is reached.
 - b. All ports are marked bad.

At each pass, the frame is created from a different data type of a palette of seven. If seven passes are requested, seven different data types are used in the test. If eight passes are requested, the first seven frames use unique data types, and the eighth is the same as the first. The seven data types are:

```

CSPAT: 0x7e, 0x7e, 0x7e, 0x7e, ...
BYTE_LFSR: 0x69, 0x01, 0x02, 0x05, ...
CHALF_SQ: 0x4a, 0x4a, 0x4a, 0x4a, ...
QUAD_NOT: 0x00, 0xff, 0x00, 0xff, ...
CQTR_SQ: 0x78, 0x78, 0x78, 0x78, ...
CRPAT: 0xbc, 0xbc, 0x23, 0x47, ...
RANDOM: 0x25, 0x7f, 0x6e, 0x9a, ....

```

Notes Do not abort this test prematurely, using **CTRL-C** or **q** to quit. Doing so may cause the test to report unexpected errors. Errors may vary depending on the hardware platform.

You cannot interrupt the test by pressing the return key (<cr>).

This command does not support High Availability (HA).

The following restrictions apply on the Brocade G620 switch:

- 32 Gbps SFPs supports 32G/16G/8G speeds
- 16 Gbps SFPs supports 16G/8G/4G speeds
- 10 Gbps SFPs supports only 10G speed
- 10 Gbps spd_mode works only with lb_mode 2 with 10 Gbps SFP only and does not work with lb_mode 1
- 16 Gbps and 32 Gbps SFPs do not support 10G speed in lb_mode 1 or 2

The Brocade 8510/DCX series cannot negotiate speeds of 1 Gbps.

This diagnostic cannot be run on an operational switch. You must disable the switch using the **chassisDisable** command before you can run this test. After the test completes, re-enable the switch using the **chassisEnable** command. Do not use the **switchDisable** command followed by manually shutting down the ports before running the test. This will not work because executing **switchEnable** after the test completes will leave the ports in the same state as before. The only way to bring the ports down before running the test and back up after the test completes is by using the **chassisDisable** command followed by the **chassisEnable** command.

The 32 Gbps-capable platforms support only lb_mode 1 and 2.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--slot slot

Specifies the number of the slot on which to run the diagnostics. All eligible ports relative to the slot number are tested. The default is 0 and designed to operate on fixed-port-count products.

-nframes count

Specifies the number of frames to send. The test progresses until the specified number of frames has been transmitted on each port. The default value is 1. The maximum number of frames sent is limited to 1000 on 16 Gbps-capable platforms and to 2000 on 32 Gbps-capable platforms. If the number of frames exceeds maximum limit, the test proceeds with the initial maximum limit of frames (for example, 1000 on 16 Gbps-capable platforms) and displays a warning message.

portLoopbackTest

-lb_mode mode

Specifies the loopback mode for the test. By default, this test uses the internal loopback. Valid values depend on the platform on which the command is executed. Unsupported values are rejected with an appropriate message. To run this test with mode values of 1 and 7, loopback cables must be connected to all front end ports or the test will fail.

1

Port Loopback (loopback plugs)

2

External (SERDES) loopback

5

Internal (parallel) loopback. This parameter is not supported in 16 Gbps-capable platforms.

7

Backend bypass & port loopback

8

Backend bypass & SERDES loopback

-spd_mode mode

Specifies the speed mode for the test. This parameter controls the speed at which each port is operated. The speed option chosen or defaulted to must not exceed the speed capability of the SFP device or the test results will be unpredictable. For example, if you use an 8 Gbps SFP in a 16 Gbps-supported slot, you must use **spd_mode 8** option to specify 8 Gbps speed. The speed modes 8, 10, 16, and 32 are supported in 32 Gbps-capable platforms.

2

Runs test at 2 Gbps.

4

Runs test at 4 Gbps.

8

Runs test at 8 Gbps (Default for 8 Gbps platforms).

10

Runs test at 10 Gbps (Default for 10 Gbps platforms). This speed mode is not supported on 16 Gbps core blades.

16

Runs test at 16 Gbps (Default for 16 Gbps platforms).

32

Runs test at 32 Gbps (Default for 32 Gbps platforms). This speed mode is not supported in 8 Gbps-capable and 16 Gbps-capable platforms.

-ports itemList

Specifies a list of blade ports to test. By default, all of the blade ports in the specified slot (**--slot**) are used. You must have root access to the **bladePortMap** command if you want to map a specific front-end blade port to a user port. You can also use the **-uports** option to specify the user ports. Refer to **itemList** for more information on the **itemList** parameter.

-uports *itemlist*

Specifies a list of user ports to test. If this operand is omitted, by default the test will run on all valid blade ports in the slot. Refer to **itemList** for more information on the *itemlist* parameter. This option is only available for 16 Gbps-capable and 32 Gbps-capable blades.

Diagnostics

When it detects failures, the test may report one or more of the following error messages. If errors persist, contact Technical Support.

DATA

Data received does not match the data sent.

ERRSTAT

Errors were found in the ASIC statistics.

INIT

Port failed to initialize.

PORTDIED

A previously initialized port went to an uninitialized state.

STATS

Errors were found in the ASIC statistics.

TIMEOUT

Did not receive a frame in the given timeout period.

XMIT

Frame transmission failure.

Examples

To run a functional test in default mode:

```
switch:admin> portloopbacktest

Running portloopbacktest .....

PASSED.
```

SEE ALSO [itemList](#)

portName

portName

Assigns or displays port names.

Synopsis

portname

portname [*slot*]/*port* [-n *name*]

portname -i [*index1*[-*index2*][...][*-f*] [-n *name*]]

portname -x [*hex1*[-*hex2*] [...][*-f*] [-n *name*]]

portname -slot *slot1*[-*slot2*] [...][*-n name*]

portname -d [*format_string* | -default | -help]

portname -h

Description

Use this command to assign a port name to a specified port or to a range of ports. The port name is included in the **portShow** output; it should not be confused with the world wide port name.

When a port name is not configured, the **portName** command displays a default name in the **portShow** output. The format of the default name is as follows:

- On standalone platforms, the default port name displays as port*portnumber*, for example, "port10."
- On enterprise-class platforms, the default port name displays as slots*slotnumber* port*portnumber*, for example, "slot1 port5."

You can identify a single port to be configured by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers (decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers. When issued without the name operand, this command displays the names of the specified ports or of all ports, if no port is specified.

Specifying multiple ports with the index (*-i* or *-x*) or slot (*-s*) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the *-i* option without a port index argument to display the **portSwap** status, or alternately use the **portSwapShow** command. When Ficon Management Server (FMS) mode is enabled, specifying a port range will fail. This is because FMS does not permit multiple ports to have the same port name. The first port in the range gets the port name, all other ports in the range will fail. Multiple ports can have the same port name in non-FMS mode.

Like all other configurable port attributes, port names persists across reboots and power cycles. They are not affected by the **configDefault** command, but they are cleared by **portCfgDefault**.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the ports to be configured, followed by a slash (/).

port

Assigns a name to a single port identified by its port number.

-i *index1*[-*index2*]

Assigns a name to a single port or to a range of ports identified by port index numbers, for example, **-i** 1/3-8 **-n** backup. You may specify multiple index ranges separated by a space, for example, **-i** 35-45 61-68 **-n** backup.

-x [*hex1* [-*hex2*]]

Specifies a port or a range of ports identified by port index numbers in hexadecimal format. You may specify multiple port ranges separated by a space, for example, **-x** 21-26 28-3c.

-f

Ignores nonexisting ports. This operand is valid only with the **-i** and **-x** options.

-slot [*slot1*[-*slot2*]]

Assigns a name to all ports on a slot or on a range of slots, for example, **-s** 3-5 **-n** backup. Multiple slot ranges are not supported with this command.

-n *name*

Specifies the name to be assigned to the ports. The port name is a character string up to 128 characters, including spaces and characters, and excluding commas (,), semicolons (;), backslashes (\), and the at sign (@). When FICON Server Management mode is enabled, the port name character string can only be up to 24 characters in length. To erase a port name, execute the port name operand as an empty string in double-quotation marks (**-n ""**). This operand is optional; if omitted, the current port name is displayed.

Some characters require a qualifier or double quotation marks when used with a bash shell; for example, enter a single quotation mark as `\'`, enter an exclamation mark as `\!`, or enter a pipe (|) as `\|`.

The empty string operand for port name (**-n ""**) will not be effective if the default configuration parameter for "Disable Default PortName" is set to **N** or enabled. Use the **configure** command to set this parameter to **Y** or disabled so that **-n ""** can be effective.

-d

Configures or displays the dynamic port name format. When executed without optional parameters, this command displays the configured dynamic port name format.

format_string

Specifies the dynamic port name format. The following control keys and their corresponding port name fields are supported:

- **S** - Switch name
- **T** - Port type
- **I** - Port index
- **C** - Slot number/port number
- **A** - Alias name
- **F** - FDMI hostname
- **R** - Remote switch name

The control keys for port name fields must be separated using periods (.),dashes (-), or underscores (_). Control keys are case-sensitive. The **format_string** must be enclosed in double quotes.

portName

- default** Sets the dynamic port name format to the default format string "S.T.I.A".
- help** Displays the syntax and usage guidelines for the **-d** operand.
- h** Displays the command usage.

Examples

To name a port tape drive 8:

```
switch:admin> portname 1/3 -n "Tape drive 8"  
switch:admin> portname 1/3  
Tape drive 8
```

To assign a name to a range of ports specified by port index numbers:

```
switch:admin> portname -i 22-26 -n backup  
switch:admin> portname -i 22-26  
port 22: backup  
port 23: backup  
port 24: backup  
port 25: backup  
port 26: backup
```

To assign a name to all ports on slot 1 and 2:

```
switch:admin> portname -s 1-2 -n backup  
switch:admin> portname -s 1-2  
port 416: backup  
port 417: backup  
port 418: backup  
port 419: backup  
port 420: backup  
port 421: backup  
(output truncated)
```

To configure dynamic port name format:

```
switch:admin> portname -d "S.T.I.R.A"
```

To display the configured dynamic port name format:

```
switch:admin> portname -d  
S.T.I.R.A
```

SEE ALSO [configDefault](#), [configure](#), [portCfgDefault](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portPeerBeacon

Sets the port peer beaconing to physically identify the interconnections between ports.

Synopsis `portpeerbeacon --enable [slot]port`

`portpeerbeacon --disable [slot]port`

`portpeerbeacon --show -all`

`portpeerbeacon --help`

Description Use this command to turn on or off the port peer beaconing to physically identify the interconnections between ports.

Using this command, the administrator can turn on the beaconing on both ends of the link and physically search the other switches or devices for the beacon pattern to find the peer port. The beaconing pattern is alternating green and amber LEDs every 1.2 seconds.

In the case of the trunk ports, the entire trunk group will be beaconed. This command is supported for more than one port or one trunk group. Newly added trunk ports are automatically included to beacon. Similarly, disabling the **portpeerbeacon** on a trunk port disables beaconing on the entire trunk group.

Port peer beaconing is supported on N_Ports, E_Ports, EX_Ports, F_Ports, AE_Ports, and trunk ports. The command is not supported on ICL ports.

Port peer beaconing is not supported if diagnostic tests are running on the port. Port peer beaconing is not supported with **spinFab** and is mutually exclusive of all the diagnostic commands.

Port peer beaconing remain unaffected when a port or trunk group is disconnected, toggled, or disabled.

The **portPeerBeacon** command is one of the commands that controls beaconing. Each command has a clearly defined scope of action:

- The **portBeacon** command enables or disables beaconing on a specified port.
- The **switchBeacon** command enables or disables beaconing on all ports in the current logical switch.
- The **chassisBeacon** command enables or disables beaconing on all ports in the chassis.
- The **portPeerBeacon** command enables or disables beaconing to identify the interconnections between ports.

The actions of the beaconing commands are independent and mutually exclusive except that the **portPeerBeacon** command overrides the **portBeacon** settings on a port, but **portBeacon** does not override port peer beaconing on a port. For example, if you enabled beaconing on the logical switch and you want to enable beaconing on the entire chassis, you must first disable switch beaconing with the **switchBeacon** command before you can use the **chassisBeacon** command to enable beaconing on the entire chassis. Likewise, existing **portBeacon** settings remain unaffected if you enable or disable beaconing on the switch or on the chassis. Failure to disable existing beaconing commands before using a different type of beaconing may cause the commands to interfere with each other in unexpected ways.

portPeerBeacon

To determine whether or not beaconing is enabled or disabled on the switch or chassis, use the **switchBeacon** or **chassisBeacon** command without operands. A value of 0 indicates that the command is disabled, a value of 1 indicates that the command is enabled. Issue the **portBeacon --show -all** command to display beaconing status. The **switchShow** command displays the status of the **switchBeacon** command only.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|--------------------|---|
| slot | Specifies the slot number on which the port peer beaconing is available. |
| port | Specifies the number of the port to be configured, relative to its slot for bladed systems. Use switchShow for a list of valid ports. |
| --enable | Enables the port peer beaconing on a particular port. |
| --disable | Disables the port peer beaconing on a particular port even when the port is disabled or offline. |
| --show -all | Displays the ports on which the Port Peer Beaconing is active. It helps to find the ports that receive the ELS and start beaconing. The administrator can verify the connectivity from the Telnet instead of physically verifying the connection. |
| --help | Displays the command usage. |

Examples To enable the port peer beacon on a port:

```
switch:admin> portpeerbeacon --enable 2/15
```

To disable the port peer beacon on a port:

```
switch:admin> portpeerbeacon --disable 15
```

To display the status of the port peer beacon on a chassis:

```
switch:admin> portpeerbeacon --show -all
PortPeerBeacon enabled port(s):
FID 128:
3/27, 3/28, 3/30, 3/31,
FID 10:
3/16, 3/17,
FID 20:
3/24, 3/25
```

To display the trunk port port peer beacon information (in this example, 2/4, 2/8, and 3/4 are the master ports of their trunking group, and the other ports are the slave ports in their trunking group):

```
switch:admin> portpeerbeacon --show -all
PortPeerBeacon enabled port(s):
FID 128:
Trunk-2/4, 2/5, 2/7, 2/0, 2/1, 2/3, 2/2, 2/6,Trunk-2/8, 2/9, 2/10, 2/11,
FID 10:
Trunk-3/4, 3/5, 3/6, 3/7,
```

SEE ALSO [chassisBeacon](#), [portBeacon](#), [switchBeacon](#)

portPerfShow

Displays port throughput performance.

Synopsis

portperfshow

```
portperfshow [[slot]port1[-[slot]port2]]
  [-tx -rx | -tx | -rx] [-t interval]
```

```
portperfshow -x hex1[-hex2]
```

```
portperfshow --help
```

Description

Use this command to display throughput information for all ports on a switch or chassis or to display the information for a specified port or port range. You can display throughput information about a single port identified by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers(decimal or hexadecimal) or by specifying a slot or a slot range.. Output includes the number of bytes received and transmitted per interval. Throughput values are displayed as either bytes, kilobytes (k), megabytes (m), or gigabytes (g). Values are rounded down.

The data is displayed one column per port plus one column that displays the total for these ports. Results display every second or over a specified interval. Press **Enter**, **Ctrl-c**, or **Ctrl-d** to terminate the command. To run this command one time only, specify an interval of zero.

When executed with the command line arguments **-tx**, **-rx**, or **-tx -rx**, this command displays the transmitter throughput, the receiver throughput, or both. For ports with status of "No_Module," "No_Light," "No_SigDet," or "Loopback" throughput is displayed as 0.

An asterisk (*) in the output indicates a SIM port that is generating or receiving traffic.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When FastWrite or Tape Pipelining is enabled, the **portPerfShow** VE link output is different. The acceleration entity (FastWrite or Tape Pipelining) responds by sending XFER_RDY and status well ahead of the actual device's response to the host. The host sends data which is stored near the device and is delivered to the device only when the device is ready. Consequently, the data may be stored near the target for some brief period of time. In this case, the **portPerfShow** output on the VE link may not match the output on the device port.

Operands

This command has the following optional operands:

```
[slot]port1[-[slot]port2]
```

Displays throughput information for a single port or for a range of ports, relative to the slot number on bladed systems. Port numbers in a range must be separated by a dash (-), for example, 3-5, or 2/0-2/15. Port ranges cannot span slots. Use **switchShow** to display a listing of valid ports. Port operands are optional; if omitted, information for all ports is displayed.

- t *time_interval***
Specifies the interval, in seconds, between each sample. The default interval is one second. If no interval is specified, the default is used. To run this command one time only, specify an interval of zero.
- tx**
Displays the transmitter throughput.
- rx**
Displays the receiver throughput.
- tx -rx**
Displays the transmitter and receiver throughput.
- x *hex1* [-*hex2*]**
Accepts an index number or a range of index numbers within the same slot in hexadecimal format as input and displays the output in slot and port number format for chassis and index number format on switches.

Examples To display performance information for all ports at a one second (default) interval:

```
switch:user> portperfshow
  0   1   2   3   4   5   6   7   8   9   10  11  12  13  14  15
=====
630.4m 0   0   0   0   0   0   0   0   0   0 630.4m 0   0   0   0   0

16  17   18   19  20  21  22  23  24  25  26  27  28  29  30  31
=====
0 210.1m 840.5m 210.1m 0   0   0   0   0   0   0   0   0   0   0   0   0

32  33  34  35  36  37  38  39  Total
=====
0   0   0   0   0   0   0   0   2.5g

  0   1   2   3   4   5   6   7   8   9   10  11  12  13  14  15
=====
630.4m 0   0   0   0   0   0   0   0   0   0 630.4m 0   0   0   0   0

16  17   18   19   20  21  22  23  24  25  26  27  28  29  30  31
=====
0 210.1m 840.6m 210.1m 0   0   0   0   0   0   0   0   0   0   0   0

32  33  34  35  36  37  38  39  Total
=====
0   0   0   0   0   0   0   0   2.5g
(output stopped)
```

To display port performance for all ports with an interval of 5 seconds:

```
switch:user> portperfshow -t 5

  0   1   2   3   4   5   6   7   8   9   10  11  12  13  14  15
=====
630.4m 0   0   0   0   0   0   0   0   0 112 630.4m 0   0   0   0   0
```

portPerfShow

```
16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
=====
0 210.1m 840.6m 210.1m 0 112 0 0 0 0 0 0 0 0 0 0
=====

32 33 34 35 36 37 38 39 Total
=====
0 0 0 0 0 0 0 0 2.5g
=====

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
=====
630.4m 0 0 0 0 0 0 0 0 0 630.4m 0 0 0 0 0
=====

16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
=====
0 210.1m 840.5m 210.1m 0 0 0 0 0 0 0 0 0 0 0 0
=====

32 33 34 35 36 37 38 39 Total
=====
0 0 0 0 0 0 0 0 2.5g
(output truncated)
```

To display performance on a single port with at a 5 second interval:

```
switch:user> portperfshow 0 -t 5
0 Total
=====
630.4m 630.4m

0 Total
=====
630.3m 630.3m
(output truncated)
```

To display transmitter throughput for a single port at a 5 second interval:

```
switch:user> portperfshow 0 -tx -t 5
0
=====
210.1m
0
=====
210.1m
(output truncated)
```

To display receiver throughput for a single port at a 5 second interval:

```
switch:user> portperfshow 0 -rx -t 5
0
=====
420.3m
0
=====
420.2m
(output truncated)
```

To display port performance on a chassis for range of ports (** indicates a SIM port):

```
switch:user> portperfshow 0-2
  0      1      2      Total
=====
*      0      *      0
```

To display port performance on a chassis for range of ports at an interval of 5 seconds:

```
switch:user> portperfshow 12/0-12/6 -t 5
      0      1      2      3      4      5      6      Total
=====
slot 12: 840.6m 0      0      0      0      0      0      630.4m  1.4g
  0      1      2      3      4      5      6      Total
=====
slot 12: 840.6m 0      0      0      0      0      0      630.4m  1.4g
  0      1      2      3      4      5      6      Total
=====
slot 12: 840.6m 0      0      0      0      0      0      630.4m  1.4g
(output truncated)
```

To display performance information for a range of ports:

```
switch:user> portperfshow -x a-b
      10      11      Total
=====
      0      0      0
```

SEE ALSO [portStatsShow](#)

portRouteShow

Displays routing tables for the specified port.

Synopsis `portrouteshow [slot/]port`

Description Use this command to display the port address ID for a specified port and the contents of the following port routing tables:

External unicast routing table

Displays how the specified port forwards unicast frames to remote domains in the following format: *domain_number: ports_bitmap*

domain_number

The remote domain ID to which frames are ultimately routed.

ports_bitmap

The port number on the ASIC pair to which frames for the domain ID forward in bitmap hex format; for example, 0x0100 indicates port 8 on the ASIC pair. The arrangement of ports on an ASIC pair is specific to the system type. For any active port, this table contains at least one entry, which routes unicast frames destined to the embedded port (value 0x10000) of the local domain.

Internal unicast routing table

Displays how the specified port forwards unicast frames to a locally attached NX_Port in the following format: *area_number: ports_bitmap*

area_number

The area number of a device (or set of looped devices) attached to the local switch.

ports_bitmap

The format of *ports_bitmap* is the same as the one used in the external unicast routing table.

Broadcast routing table

Displays how the specified port forwards broadcast frames. There is one bit map entry in this table, similar to the bit maps in the other tables; however, this table typically has only Bit 16 set (value 0x10000), indicating this port always routes broadcast frames to the embedded port, for handling by the firmware.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to display, followed by a slash (/).

port

Specifies the number of the port to display, relative to its slot for bladed systems.

Examples To display the routing tables for a port:

```
switch:user> portrouteshow 4/15
port address ID: 0x02bf00
external unicast routing table:
  1: 0x4 (vc=3)
  2: 0x10000 (vc=0)

internal unicast routing table:
  60: 0x8000 (vc=2)
  63: 0x1000 (vc=5)
broadcast routing table:
  0x10000
```

SEE ALSO [bcastShow](#), [fabricShow](#), [switchShow](#), [topologyShow](#), [uRouteShow](#)

portShow

Displays status and configuration parameters for ports and GbE ports.

Synopsis

```
portshow [slot][ge]port
```

```
portshow [slot]port -link [-force]
```

```
portshow -i [ index1 [-index2] [...] ] [-f]
```

```
portshow -pid pid
```

```
portshow -x [hex1[-hex2] [...]]
```

```
portshow option [slot]ge_port [optional_args] [--filter filter_args]
```

```
portshow option [all | ve_port] arguments  
[optional_arguments] [--filter filter_args]
```

```
portshow option [all | ge_port] arguments  
[optional_arguments] [--filter filter_args]
```

Description

Use this command to display general port status and specific configuration parameters for a specified port, GbE port, or VE_Port.

If this command is executed for a specified port with no additional options, it displays general status and configuration for that port. If executed with optional arguments for a Gigabit Ethernet (GbE) port or VE_Port, the command displays extension-related port configuration parameters specific to the Brocade 7840 switch, and the Brocade FX8-24 and the Brocade SX6 extension blades.

The behavior of this command is platform-specific. Some command options are not available on all platforms. Use the following section headings to navigate this page.

- Display general port status information on all platforms
- Display IP Interface configurations on the Brocade 7840, SX6, or FX8-24 platforms:
 - **portshow ipif** - Displays the local IP interfaces.
 - **portshow arp** - Displays the content of the address resolution protocol (ARP) table.
 - **portshow iproute** - Displays static routes on the IP interface.
 - **portshow vlantag** - Displays the IP interface VLAN configuration. Supported on the Brocade FX8-24 only.
 - **portshow autoneg** - Displays autonegotiation status.
 - **portshow ipsec-policy** - Displays the IPSEC Policy information. Supported on Brocade 7840 switch and the Brocade SX6 extension blade only.
 - **portshow mgmtif** - Displays the inband management interfaces. Supported on the Brocade FX8-24 only.
 - **portshow mgmtroute** - Displays the routes for the inband management interfaces. Supported on the Brocade FX8-24 only.
 - **portshow lag** - Displays the static link aggregation group (LAG) information. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
 - **portshow tcl** - Displays the Traffic Control List (TCL) information. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
 - **portshow lan-stats** - Displays the LAN statistics. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
 - **portshow sla** - Displays the Service Level Agreement (SLA) statistics. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- Display extension tunnels, circuits, management interfaces and FICON statistics on the Brocade 7840, SX6, or FX8-24 platforms:
 - **portshow fcipunnel** - Displays extension tunnels.
 - **portshow fcipcircuit** - Displays extension circuits.
 - **portShow xtun** - Displays FICON and FCP emulation statistics and current runtime conditions.

To display the command usage on the switch, use **portShow [action]**.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Some of the features supported by this command may require a license.

In an AD context, if one of the L_Ports or NPIV ports is a part of the current AD, the complete device information attached to the port is shown in the output.

Function General port status display commands supported on all platforms

Synopsis `portshow [slot][ge]port`

`portshow -i [index1[-index2] [...]] [-f]`

`portshow -x [hex1[-hex2] [...]]`

`portshow -pid pid`

Description Use this command to display general port status and configuration parameters for the specified port. This command is valid on all platforms, but the output is platform-specific and not all fields are displayed on all platforms.

The following general information is displayed when the command is issued for a non-GbE port without additional arguments:

| | |
|--------------------------|---|
| portIndex | Index number assigned to the port. |
| portName | Name assigned to the port by the portName command. On standalone platforms, the default port name is the port number, for example, "port5". On enterprise-class platforms, the default port name is a combination of the slot number and the port number, for example, "slot1port5". |
| portHealth | Current health of the port (requires a Fabric Vision license). |
| Authentication | Authentication type and associated parameters (if applicable) used on the port at port online. |
| None | No authentication was performed. |
| FCAP | FCAP authentication was performed. |
| DHCHAP | DHCHAP authentication was performed. Also displays DH group and hash type used for authentication. |
| portDisableReason | Provides an explanation for the port's disabled status. |
| portCFlags | Port control flags. |
| portFlags | A bit map of port status flags, including information on the type of port, whether it is fully online, and whether logins have been accepted. The port flags display ENCRYPT if the port has been enabled for encryption. The port flags display COMPRESS if the port has been enabled for compression. The port flags display D_PORT if the port has been enabled as a diagnostic port. The port flags display FLOGI_LOGO if the base device logs out and at least one NPIV is online. |
| portType | The port's type and revision numbers. |
| POD Port | Ports on Demand License status. |
| portState | The port's SNMP state: |

| | |
|------------------------------|---|
| Online | Up and running. |
| Offline | Not online, see portPhys for more detail. |
| Testing | Running diagnostics. |
| Faulty | Failed diagnostics. |
| Persistently Disabled | Persistently disabled. |
| Protocol | Protocol used by the port: FC or FCoE. |
| portPhys | The port's physical state: |
| No_Card | No interface card present. |
| No_Module | No module (GBIC or other) present. |
| No_Light | Module is not receiving light (8 Gbps-capable ports only). |
| No_SigDet | No signal detected; displays when a quad small form-factor pluggables (QSFPs) is installed without cables). |
| Mod_Inv | Incompatible vendor or module speed mismatch. |
| No_Sync | Receiving light but out of sync. |
| In_Sync | Receiving light and in sync. |
| Laser_Flt | Module is signaling a laser fault. |
| Port_Flt | Port marked faulty. |
| Diag_Flt | Port failed diagnostics. |
| Lock_Ref | Locking to the reference signal. |
| portScn | The port's last State Change Notification. |

| | |
|--|---|
| port generation number | The port's generation number for the last offline state change. |
| portId | The port's 24-bit port ID. |
| portIfId | The user port's interface ID. |
| portWwn | The port's world wide name. |
| portWwn of devices(s) connected | The World Wide Port Names of connected devices. If the base device logs out and one or more NPIVs are online, the PWWN of the NPIVs online are displayed. |
| Distance | The port's long-distance level. In the case of LD mode, the user configured distance and actual distance also are displayed. See portCfgLongDistance for information on long distance levels. |
| Port part of other AD | Yes or No. |
| portSpeed | The port's fixed speed (1, 2, 4, or 8 Gbps) or negotiated speed (N1 Gbps, N2 Gbps, N4 Gbps, N8 Gbps or AN). |
| FEC | Forward error correction (FEC) status: displays "active" if FEC is enabled on the port and the port is online; displays "inactive" if the feature is disabled. Refer to the portCfgFec command for more information. |
| LE domain | The LE domain ID. |
| FC FastWrite | The status of FC FastWrite (ON or OFF) (Deprecated). |
| If the port is configured as an EX_Port, the following additional port information is displayed: | |
| EX_Port Mode | The port is configured as an EX_Port. |
| Fabric ID | The fabric ID assigned to this EX_Port; this is the fabric ID of the edge fabric attached to this EX_Port. |
| Front Phantom | Information on the front phantom domain presented by this EX_Port. Includes the preferred (if not active) or actual (if active) domain ID for the front domain and the WWN of the front domain. |
| Pr Switch Info | Information on the principal switch of the edge fabric attached to this EX_Port. Includes the domain ID and WWN of the principal switch. |

| | |
|--|---|
| BB XLate | Information on the xlate (translate) phantom domain presented at this port. Includes the preferred (if not active) or actual (if active) domain ID for the xlate phantom domain and the WWN of the xlate phantom domain. The xlate phantom domain connected at this port is in the same fabric as the router and represents the edge fabric connected to the EX_Port. |
| Authentication Type | Displays NONE or DH-CHAP. DH-CHAP is the only authentication type supported on EX_Ports. |
| DH Group | Displays DH group [0-4] if DH-CHAP authentication is used. Otherwise displays N/A. |
| Hash Algorithm | Displays hash type (MD5 or SHA-1) if DH-CHAP authentication is used. Otherwise, displays N/A. |
| Edge fabric's primary WWN | If the EX_Port is connected to an edge switch with FCS policy enforcement, the WWN of the primary FCS is displayed when the edge fabric is secure and the primary FCS is online. Otherwise, displays "No Primary". |
| Edge fabric's version stamp | If the EX_PORT is connected to an edge switch with FCS policy enforcement, the version of the security database is displayed. Otherwise displays N/A. |
| <p>The portShow command displays FCoE ports with "Protocol: FCoE" and "portSpeed: 10Gbps". Only a subset of information is displayed. Refer to the Examples section for an illustration. Use fcoe --cfgshow and fcoe --loginshow to display FCoE-specific configuration details.</p> <p>Following the general information, the command displays three columns of counters. The first column shows interrupt statistics:</p> | |
| Interrupts | Total number of interrupts. |
| Unknown | Interrupts that are not counted elsewhere. |
| Lli | Low-level interface (physical state, primitive sequences). |
| Proc_rqrd | Frames delivered for embedded N_Port processing. |
| Timed_out | Frames that have timed out. |
| Rx_flushed | Frames requiring translation. |
| Tx_unavail | Frames returned from an unavailable transmitter. |
| Free_buffer | Free buffer available interrupts. |

| | |
|-------------------------------|---------------------------------------|
| Overrun | Buffer overrun interrupts. |
| Suspended | Transmission suspended interrupts. |
| Parity_err | Central memory parity errors. |
| 2_parity_err | Secondary transmission parity errors. |
| CMI_bus_err | Control message interface errors. |
| FEC Corrected Blocks | FEC corrected block errors. |
| FEC Uncorrected Blocks | FEC uncorrected block errors. |

The second column displays link error status block counters.

The third column shows the number of F_RJTs and F_BSYs generated. For L_Ports, the third column also displays the number of loop initialization protocols (LIPs) received, number of LIPs transmitted, and the last LIP received.

Operands

This command has the following port operands:

| | |
|---|---|
| <i>slot</i> | For bladed systems only, specifies the slot number of the port to be displayed, followed by a slash (/). |
| <i>port</i> | Specifies the number of the port to be displayed, relative to its slot for bladed systems. Use switchShow for a listing of valid port numbers. |
| -link [-force] | Displays the diagnostic information from cached data for a local switch port and the peer port. The -force option displays the real-time diagnostic information from the registers for a port. |
| -i <i>index1</i>[-<i>index2</i>] | Specifies a port or a range of ports identified by port index numbers. You can specify multiple index ranges separated by a space, for example, 33-38 40-60. Port indexes are supported only if PortSwap is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the portSwap status, or alternately use portSwapShow . |
| -x [<i>hex1</i> [-<i>hex2</i>]] | Specifies a port or a range of ports identified by port numbers, index number in hexadecimal format. You may specify multiple port ranges separated by a space, for example, -x 21-26 28-3c. |
| -pid <i>pid</i> | Specifies the diagnostics information of a remote switch port along with its peer port. |

Examples To display the current state of a D_Port:

```
switch:admin> portshow 28
portName: port28
portHealth: Fabric vision license not present

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x10004103    PRESENT ACTIVE E_PORT \
    G_PORT D_PORT U_PORT LOGIN LED
LocalSwcFlags: 0x0
portType: 24.0
POD Port: Port is licensed
portState: 1    Online
Protocol: FC
portPhys: 6 In_Sync    portScn: 0    Flow control mode 4
port generation number:    0
state transition count:    1

portId:    2d1c00
portIfId:    43020028
portWwn:    20:1c:00:05:33:13:2f:b3
portWwn of device(s) connected:

Distance:    normal
portSpeed:    8Gbps

FEC: Inactive
LE domain: 0
FC Fastwrite: OFF
Interrupts:    0    Link_failure: 0    Frjt:    0
Unknown:    0    Loss_of_sync: 1    Fbsy:    0
Lli:    14    Loss_of_sig: 2
Proc_rqrd:    4    Protocol_err: 0
Timed_out:    0    Invalid_word: 0
Rx_flushed:    0    Invalid_crc: 0
Tx_unavail:    0    Delim_err: 0
Free_buffer:    0    Address_err: 0
Overrun:    0    Lr_in:    2
Suspended:    0    Lr_out:    0
Parity_err:    0    Ols_in:    0
2_parity_err:    0    Ols_out:    1
CMI_bus_err:    0
```

To display the current state of a port with encryption enabled:

```
switch:admin> portshow 10/44
portIndex: 348
portName: slot10port44
portHealth: Fabric vision license not present

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x10000103    PRESENT ACTIVE E_PORT T_PORT\
    T_MASTER G_PORT U_PORT ENCRYPT LOGIN
LocalSwcFlags: 0x0
portType: 24.0
```

portShow

```
portState: 1   Online
Protocol: FC
portPhys: 6   In_Sync   portScn: 1   Online Trunk master port
port generation number:    44
state transition count:    12
```

To display the state of a VE_Port on a chassis with an FX8-24 blade:

```
switch:admin> portshow 7/12
portIndex: 140
portName: slot7port12
portHealth: Fabric vision license not present

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x490b PRESENT ACTIVE VIRTUAL E_PORT \
    G_PORT U_PORT LOGICAL_ONLINE LOGIN LED
portType: 12.0
portState: 1   Online
Protocol: FC
portPhys: 255 N/A   portScn: 16   E_Port
port generation number:    1048
state transition count:    26

portId:    018c00
portIfId:  43720806
portWwn:   20:8c:00:05:1e:7a:7a:00
portWwn of device(s) connected:

Distance:  normal
```

To display port status for a GbE port on a Brocade 7840:

```
switch:admin> portshow ge2
Eth Mac Address: 00.05.33.65.82.ca
Port State: 1   Online
Port Phys: 6   In_Sync
Port Flags: 0x3 PRESENT ACTIVE
Port Speed: 10G
```

To display an FCoE port:

```
switch:admin> portshow 8
portIndex: 8
portName:
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x2cb03 PRESENT ACTIVE F_PORT G_PORT U_PORT\
    LOGICAL_ONLINE LOGIN NOELP LED NSREG ACCEPT FLOGI
portType: 17.0
POD Port: Port is licensed
portState: 1   Online
Protocol: FCoE
portPhys: 6   In_Sync   portScn: 32   F_Port
port generation number:    0
```

```

state transition count:    1

portId:    850800
portIfId:    43020028
portWwn:    20:08:00:05:1e:76:60:80
portWwn of device(s) connected:

Distance:    normal
portSpeed:    10Gbps

```

To display the cached data for a port:

```

switch:admin> portshow 32
portIndex:    32
portName:    port32
portHealth:
Authentication:    None
portDisableReason:    None
portCFlags:    0x1
portFlags:    0x18024b03 PRESENT ACTIVE T_FPORT T_FMASTER F_PORT G_PORT U_PORT
LOGICAL_ONLINE LOGIN NOELP LED ACCEPT FLOGI
LocalSwcFlags:    0x0
portType:    24.0
POD Port:    Port is licensed
portState:    1Online
Protocol:    FC
portPhys:    6In_Sync portScn:    32F_Port    Trunk master port
port generation number:    22
state transition count:    1
portId:    323600
portIfId:    43020027
portWwn:    20:20:00:27:f8:81:85:69
Logical portWwn:    50:02:7f:88:18:58:60:32
portWwn of device(s) connected:
25:00:00:27:f8:65:02:3e
Distance:    normal
portSpeed:    N8Gbps
FEC:    Inactive
Credit Recovery:    Inactive
Aq:    Inactive
FAA:    Inactive
F_Trunk:    Active
LE domain:    0
FC Fastwrite:    OFF
Interrupts:
Unknown:
Lli:
Proc_rqrd:
Timed_out:
Rx_flushed:
Tx_unavail:
Free_buffer:
Overrun:
Suspended:
Parity_err:
2_parity_err:
CMI_bus_err:
nodeWwn:
PN_PORT Phy Type:
FEC Corrected Blocks:

```

| | | | | | |
|---------------|-------|---------------|---|-------|---|
| Interrupts: | 0 | Link_failure: | 0 | Frjt: | 0 |
| Unknown: | 0 | Loss_of_sync: | 0 | Fbsy: | 0 |
| Lli: | 28 | Loss_of_sig: | 0 | | |
| Proc_rqrd: | 24868 | Protocol_err: | 0 | | |
| Timed_out: | 0 | Invalid_word: | 0 | | |
| Rx_flushed: | 0 | Invalid_crc: | 0 | | |
| Tx_unavail: | 0 | Delim_err: | 0 | | |
| Free_buffer: | 0 | Address_err: | 0 | | |
| Overrun: | 0 | Lr_in: | 4 | | |
| Suspended: | 0 | Lr_out: | 2 | | |
| Parity_err: | 0 | Ols_in: | 1 | | |
| 2_parity_err: | 0 | Ols_out: | 4 | | |
| CMI_bus_err: | 0 | | | | |

portShow

```
FEC Uncorrected Blocks: 0
```

```
PEER PORT
```

```
portWwn: 20:20:00:27:f8:81:85:69
nodeWwn: 20:20:00:27:f8:81:85:69
PN_PORT Phy Type: PN Port/PF Port [PN Port/PF Port, lossless Ethernet MAC]
Link_failure: 0          Loss_of_sync: 0          Loss_of_sig: 0
Protocol_err: 0          Invalid_word: 0          Invalid_crc: 0
FEC Corrected Blocks: 0
FEC Uncorrected Blocks: 0
```

To display the information of a remote switch port:

```
switch:admin> portshow -pid 0x010500
SWITCH PORT0xDDAA00
```

```
portWwn: 20:20:00:27:f8:81:85:69
nodeWwn: 20:20:00:27:f8:81:85:69
PN_PORT Phy Type: PN Port/PF Port
Link_failure: 0          Loss_of_sync: 0          Loss_of_sig: 0
Protocol_err: 0          Invalid_word: 0          Invalid_crc: 0
```

```
PEER PORT 0xDDAAxx
```

```
portWwn: 20:20:00:27:f8:81:85:69
nodeWwn: 20:20:00:27:f8:81:85:69
PN_PORT Phy Type: PN Port/PF Port
Link_failure: 0          Loss_of_sync: 0          Loss_of_sig: 0
Protocol_err: 0          Invalid_word: 0          Invalid_crc: 0
```

Function Display IP Interface configurations on the Brocade 7840, SX6, or FX8-24 platforms

Synopsis `portshow option [slot]ge_port [optional_args] [--filter filter_args]`

Description Use this command to display FCIP-related configuration parameters on the Brocade 7840 switch, on the Brocade SX6 extension blade, and on the Brocade FX8-24 blade. The parameters displayed by this command are set with the **portCfg** command. The following displays are supported with this command:

- **portshow ipif** - Displays the local IP interfaces.
- **portshow iproute** - Displays static routes on the IP interface.
- **portshow arp** - Displays the content of the address resolution protocol (ARP) table.
- **portshow vlantag** - Displays the IP interface VLAN configuration. Supported on the Brocade FX8-24 only.
- **portshow autoneg** - Displays autonegotiation status.
- **portshow ipsec-policy** - Displays IPSEC Policy information. Supported on Brocade 7840 switch and the Brocade SX6 extension blade only.
- **portshow mgmtif** - Displays the inband management interfaces.
- **portshow mgmtroute** - Displays the routes for the inband management interfaces.
- **portshow lag** - Displays the static link aggregation group (LAG) information. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- **portshow tcl** - Displays the Traffic Control List (TCL) information. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- **portshow lan-stats** - Displays the LAN statistics. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.
- **portshow sla** - Displays the SLA statistics. Supported on the Brocade 7840 switch and the Brocade SX6 extension blade only.

Notes IPv6 addresses are supported.

Operands This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be displayed, followed by a slash (/).

ge_port

Specifies the number of the GbE port to be displayed relative to the slot number. The GbE ports are numbered ge0 - ge9 on the Brocade FX8-24 blade. The 10GbE ports on the Brocade FX8-24 blade are numbered xge0 and xge1. Brocade 7840 switch and the Brocade SX6 extension blade have two 40GbE ports labeled ge0 and ge1, and 16 1 or 10GbE ports labeled ge2-ge17. Use the **switchShow** command for a listing of valid ports

ipif

Displays the IP interface ID, IP address, netmask, and MTU for IPv4 addresses. Displays the prefix instead of the netmask for IPv6 addresses. Flags are explained in the command output (refer to the example below).

--link-local

Displays the link-local interface addresses.

iproute

Displays the IP address, netmask, gateway, metrics, and flags. Displays the prefix instead of the netmask for IPv6 addresses.

| | |
|---------------------|---|
| --link-local | Displays the link-local interface routes. This operand is optional. |
| arp | Displays the address resolution protocol (ARP) and IPv6 neighbor discovery table. You can display the content of the ARP table, but you cannot modify its contents. |
| --link-local | Displays the neighbor discovery entries for the link-local addresses. This operand is optional. |
| -lmac | Displays the local MAC address. This operand is optional. |
| vlan tag | Displays the VLAN Tagging configuration. For each entry, the output displays the IP interface address, the destination IP address, the VLAN ID, the L2 CoS priority, and a flag. This display includes tunnel- and IPIF-level configurations. The following flags indicate the type of configuration: |
| Perm | Permanent entry. Permanent entries are configured at the IP interface level with the portCfg vlan tag command. |
| Perm Net | Network-wide permanent entry. The same as the permanent entry except that no destination address was defined (defaults to 0.0.0.0), so traffic to all destinations is tagged. |
| App | Application layer VLAN configuration defined at the circuit level. |
| autoneg | Displays the auto-negotiation status. |
| ipsec-policy | Displays IPSEC Policy information. Currently supported on Brocade 7840 switch only. |
| mgmtif | Displays the inband management interfaces configured for a specified GbE Port or for all GbE Ports. For each GbE Port, the display includes the interface status (enabled or disabled), the interface IPv4 Address, the netmask, effective MTU, and annotated port flags. |
| mgmtroute | Displays the management routes configured for a specified GbE Port or for all GbE Ports. For each GbE Port, the display includes the destination IPv4 Address, the netmask, the gateway address, and annotated port flags. |
| lag | Displays the static LAG information. For each entry, the output displays the LAG name, state, port count, and member ports. |
| --detail | Displays a full view configuration details for all LAGs. The output displays the operational state of LAG, port count, and the admin state, operational state, speed, and autonegotiation status for the port members. |

tcl

Displays TCL information. The following options are supported:

- s | --summary**
Displays summary view of all the TCLs.
- d | --detail**
Displays detailed view of all the TCLs.
- p | --priority**
Sorts the TCL list based on the priority.
- s | --sort *sort_field***
Sorts the TCL list based on specified sort field. Valid values for *sort_field* are **name** (TCL name), **priority** (priority ID of the TCL), **src-addr** (source IP address), and **dst-addr** (destination IP address).

portshow lan-stats *actions [args]*

Displays the LAN IP extension statistics.

- per-flow**
Displays the details of the LAN connections based on the throughput. If the number of connections is more than 25, the top 25 throughput connections are displayed. If the number of connections is less than 25, all the available connections are displayed.
- top *count***
Displays the specified number of top throughput connections.
- bottom *count***
Displays the specified number of bottom throughput connections.
- tcp**
Displays the TCP error counter values for the LAN connections. Use this operand with **-top** or **-bottom** to display TCP error counter values for the specified number of top or bottom throughput connections.
- application [*app_name* | *port*]**
Displays aggregate throughput of the well-known applications.
 - app_name***
Displays the aggregation throughput and individual connection information which contributes towards the aggregated throughput for the specified application.
 - port***
Displays the aggregation throughput and individual connection information which contributes towards the aggregated throughput for the specified port number.
- per-flow -compression**
Displays the details of the LAN connections based on compression. If the number of connections is more than 25, the top 25 throughput connections are displayed. If the number of connections is less than 25, all the available connections are displayed.
- top *count***
Displays the specified number of top compression connections.
- bottom *count***
Displays the specified number of bottom compression connections.

| | |
|-------------------------|---|
| -tcp | Displays the TCP error counter values for the LAN connections. Use this operand with -top or -bottom to display TCP error counter values for the specified number of top or bottom compression connections. |
| --known-apps | Displays the list of supported well-known applications for the aggregation statistics. |
| --global | Displays the global LAN statistics for the switch. |
| -dp | Displays the global LAN statistics of each dual processor within the switch. |
| -reset | Resets the global LAN statistics. |
| -lifetime | Displays the entire lifetime global LAN statistics even after the statistics are reset. |
| --hist-stats | Displays the most recently closed per-flow LAN statistics. |
| -all | Displays a summary view of the per-flow LAN statistics. The following operand is required: |
| -index | Specifies a port identified by port index number. |
| -detail | Displays details about the LAN statistics including TCP flags used to close the connection. The following operands are required: |
| -dp | Displays the global LAN statistics of each dual processor within the switch. |
| -index | Specifies a port identified by port index number. |
| -freeze | Freezes the LAN connection statistic table. |
| -thaw | Clears the static LAN connection statistic table. |
| sla [name all] | Displays the configuration parameters for a specific SLA or all SLAs. The following optional arguments are supported: |
| -s --summary | Displays a summary view of the configuration parameters for a specific SLA or all SLAs. |
| -d --detail | Displays a detailed view of the configuration parameters for a specific SLA or all SLAs. |

--filter filter_args

Filters the **portShow** output based on the specified filter arguments. The filter arguments can be specified in any combination or as a conditional statement using the logical AND or OR operator. A conditional statement can have up to 30 conditions. This operand is supported with **ipif**, **iproute**, **tcl**, and **lan-stats** options. The following filter arguments are supported:

-ipaddr ip_address

Filters the output based on the specified IP address.

-port [slot]port

Filters the output based on the specified port number.

-tcp-port port

Filters the output based on the specified TCP port number.

-dp [slot]dp#

Filters the output based on the specified dual processor ID.

-retransmits value

Filters the output based on the retransmits exceeding specified value.

-rtt ms

Filters the output based on the specified circuit round trip time in milliseconds.

-connect-count value

Filters the output based on the connected count exceeding specified value.

-bps value[k | m]

Filters the output based on bandwidth (bytes per second) exceeding the specified value. Specify **k** for KBps and **m** for MBps.

-or

The logical OR operator.

-and

The logical AND operator.

-default show | hide

Sets the default display action if the specified filter statement is not supported. The default action is **hide**.

Examples

To display the IP interface and static route configured for failover crossports on the Brocade FX8-24:

```
switch:admin> portshow ipif 8/xge0
Port          IP Address          / Pfx  MTU   VLAN  Flags
-----
8/xge0        192.168.10.20       / 24   1500  n/a   U R M
8/xge0        192.168.11.21       / 24   1500  n/a   U R M X
-----
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
       N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

switch:admin> portshow iproute 8/xge0
Port          IP Address          / Pfx  Gateway  Flags
-----
8/xge0        192.168.0.0         / 24   *         U C
```

portShow

```

8/xge0      192.168.0.111      / 32  *      U H L
8/xge0      192.168.1.0        / 24  *      U C X
8/xge0      192.168.1.111     / 32  *      U H L X

```

Flags: U=Usable G=Gateway H=Host C=Created(Interface)
S=Static L=LinkLayer X=Crossport

To display the IP interface and static route configured on Brocade 7840:

```

switch:admin> portshow iproute -l
Port          IP Address          / Pfx  Gateway          Flags
-----
ge0.dp0       192.168.0.0         / 24  *                U C
ge0.dp0       192.168.0.1         / 32  *                U H L
ge0.dp0       192.168.2.0         / 24  192.168.0.1     U G S
ge0.dp0       fe80::              / 64  *                U C
ge0.dp0       fe80::5:33ff:f065:7b08 / 128 *                U H L
ge0.dp0       ff01::              / 32  *                U C
ge0.dp0       ff02::              / 32  *                U C
ge0.dp1       192.168.0.0         / 24  *                U C
ge0.dp1       fe80::              / 64  *                U C
ge0.dp1       fe80::5:33ff:f165:7b08 / 128 *                U H L
ge0.dp1       ff01::              / 32  *                U C
ge0.dp1       ff02::              / 32  *                U C
ge1.dp0       192.168.1.0         / 24  *                U C
...
ge17.dp1      fe80:11::5:33ff:f165:7b19 / 128 *                U H L
ge17.dp1      ff01:11::           / 32  *                U C
ge17.dp1      ff02:11::           / 32  *                U C

```

Flags: U=Usable G=Gateway H=Host C=Created(Interface)
S=Static L=LinkLayer X=Crossport

To display the IP interfaces configured on Brocade 7840:

```

switch:admin> portshow ipif -l
Port          IP Address          / Pfx  MTU   VLAN  Flags
-----
ge0.dp0       fe80::5:33ff:f065:7b08 / 64  1500  0     U R M
ge0.dp0       192.168.0.10         / 24  1500  0     U R M
ge0.dp1       fe80::5:33ff:f165:7b08 / 64  1500  0     U R M
ge0.dp1       192.168.0.11         / 24  1500  0     U R M
ge1.dp0       fe80:1::5:33ff:f065:7b09 / 64  1500  0     U R M
ge1.dp0       192.168.1.10         / 24  1236  100   U R M
ge1.dp1       fe80:1::5:33ff:f165:7b09 / 64  1500  0     U R M
ge1.dp1       2000::10             / 64  1500  0     U R M
ge2.dp0       fe80:2::5:33ff:f065:7b0a / 64  1500  0     U R M
ge2.dp1       fe80:2::5:33ff:f165:7b0a / 64  1500  0     U R M
ge3.dp0       fe80:3::5:33ff:f065:7b0b / 64  1500  0     U R M
ge3.dp1       fe80:3::5:33ff:f165:7b0b / 64  1500  0     U R M
ge4.dp0       fe80:4::5:33ff:f065:7b0c / 64  1500  0     U R M
ge4.dp1       fe80:4::5:33ff:f165:7b0c / 64  1500  0     U R M
...
ge11.dp1      fe80:b::5:33ff:f165:7b13 / 64  1500  0     U R M
ge12.dp0      fe80:c::5:33ff:f065:7b14 / 64  1500  0     U R M
ge12.dp1      fe80:c::5:33ff:f165:7b14 / 64  1500  0     U R M
ge13.dp0      fe80:d::5:33ff:f065:7b15 / 64  1500  0     U R M
ge13.dp1      fe80:d::5:33ff:f165:7b15 / 64  1500  0     U R M

```

```

ge14.dp0    fe80:e::5:33ff:f065:7b16    / 64  1500  0    U R M
ge14.dp1    fe80:e::5:33ff:f165:7b16    / 64  1500  0    U R M
ge15.dp0    fe80:f::5:33ff:f065:7b17    / 64  1500  0    U R M
ge15.dp1    fe80:f::5:33ff:f165:7b17    / 64  1500  0    U R M
ge16.dp0    fe80:10::5:33ff:f065:7b18    / 64  1500  0    U R M
ge16.dp1    fe80:10::5:33ff:f165:7b18    / 64  1500  0    U R M
ge17.dp0    fe80:11::5:33ff:f065:7b19    / 64  1500  0    U R M
ge17.dp1    fe80:11::5:33ff:f165:7b19    / 64  1500  0    U R M

```

```

-----
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
      N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

```

To display the ARP tables on the Brocade 7840:

```
switch:admin> portshow arp
```

| Port | IP Address | MAC Address | Flags |
|---------|--------------|-------------------|------------------|
| ge6.dp0 | 77.195.6.1 | 00:00:00:00:00:00 | Dynamic |
| ge6.dp0 | 192.168.0.10 | 00:05:33:65:84:0e | Dynamic Resolved |
| ge8.dp0 | 10.1.8.76 | 00:05:33:65:84:10 | Dynamic Resolved |

```
switch:admin> portshow arp ge6
```

| Port | IP Address | MAC Address | Flags |
|---------|--------------|-------------------|------------------|
| ge6.dp0 | 77.195.6.1 | 00:00:00:00:00:00 | Dynamic |
| ge6.dp0 | 192.168.0.10 | 00:05:33:65:84:0e | Dynamic Resolved |

To display the auto-negotiation of 1GbE port:

```

switch:admin> portshow autoneg 1/ge2
Auto-Negotiation Configuration:
Port: 1/ge2
Auto-Negotiate: Enabled

```

To display the IPSEC policy for a Brocade 7840 switch:

```
switch:admin> portshow portshow ipsec-policy --ike
```

| IPSec Policy | Key | IKE-ID | Oper | Flg | Local-Addr | Remote-Addr | IKE Rekey | ESP Rekey |
|--------------|--------------|--------|------|-----|--------------|---------------|-----------|-----------|
| policy1 | 123456789012 | | | | | | | |
| dp0.0 | UP | R | | | 192.168.0.20 | 192.168.0.120 | - | - |
| dp0.1 | UP | I | | | 192.168.0.20 | 192.168.0.121 | 2h44m4s | 0 |
| dp1.0 | UP | I | | | 192.168.0.21 | 192.168.0.120 | 2h43m58s | 0 |

```
Flags: *=Name Truncated I=Initiator R=Responder
```

To display the LAG information:

```
switch:admin> portshow lag
LAG Name      State      Port Count  Member Ports
-----
lag0          Online     3           ge2,ge3,ge4
edgeSw1       Online     2           ge7,ge8
-----
```

To display detailed LAG information:

```
switch:admin> portshow lag --detail
LAG: lag0
-----
Oper State:   Online
Port Count:   3
Port  AdminSt  OperState  Speed  AutoNeg
ge2    Enabled   Online     10G    Enabled
ge3    Enabled   Online     10G    Enabled
ge4    Enabled   Online     10G    Enabled

LAG: edgeSw1
-----
Oper State:   Online
Port Count:   2
Port  AdminSt  OperState  Speed  AutoNeg
ge7    Enabled   Online     1G     Disabled
ge8    Enabled   Online     1G     Disabled
```

To display detailed output of all the TCLs:

```
switch:admin> portshow tcl -d
TCL: hostAtoB
=====
Admin Status:      Enabled
Priority:           10
Target:            24-Medium (tid:9)
VLAN:              ANY
L2COS:             ANY
DSCP:              ANY
Source Address:    10.0.0.0/8
Destination Address: ANY
L4 Protocol:       ANY
Protocol Port:     ANY
Segment Preservation: Disabled
Action:            Allow
Cfgmask:           0x085c3a27
Hit Count:         0

TCL: default
=====
Admin Status:      Enabled
Priority:           65535
Target:            -
VLAN:              ANY
L2COS:             ANY
DSCP:              ANY
Source Address:    ANY
Destination Address: ANY
```

```

L4 Protocol:      ANY
Protocol Port:    ANY
Segment Preservation: Disabled
Action:           Deny
Cfgmask:          0x08c9007
Hit Count:        0

```

To display the LAN connections based on the throughput (in the following example, the number of connections is more than 25):

```

switch:admin> portshow lan-stats --per-flow
***Displaying Top 25 connections by throughput***

```

| Src-Address | Dst-Address | Sport | Dport | Pro | Tx(B/s) | Rx(B/s) |
|----------------|----------------|-------|-------|-----|---------|---------|
| 192.168.20.223 | 192.168.10.225 | 49678 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49695 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49672 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49687 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49690 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49689 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49693 | 49864 | TCP | 10.5m | 10.5m |
| 192.168.20.223 | 192.168.10.225 | 49683 | 49864 | TCP | 10.4m | 10.4m |
| 192.168.20.223 | 192.168.10.225 | 49679 | 49864 | TCP | 10.4m | 10.4m |
| 192.168.20.223 | 192.168.10.225 | 49674 | 49864 | TCP | 10.4m | 10.4m |
| 192.168.20.223 | 192.168.10.225 | 49694 | 49864 | TCP | 10.4m | 10.4m |
| 192.168.20.223 | 192.168.10.225 | 49671 | 49864 | TCP | 10.4m | 10.4m |
| 192.168.20.223 | 192.168.10.225 | 49691 | 49864 | TCP | 10.4m | 10.4m |
| 192.168.20.223 | 192.168.10.225 | 49670 | 49864 | TCP | 10.3m | 10.3m |
| 192.168.20.223 | 192.168.10.225 | 49675 | 49864 | TCP | 10.3m | 10.3m |
| 192.168.20.223 | 192.168.10.225 | 49680 | 49864 | TCP | 10.3m | 10.3m |
| 192.168.20.223 | 192.168.10.225 | 49676 | 49864 | TCP | 10.3m | 10.3m |
| 192.168.20.223 | 192.168.10.225 | 49686 | 49864 | TCP | 10.3m | 10.3m |
| 192.168.20.223 | 192.168.10.225 | 49673 | 49864 | TCP | 10.0m | 10.0m |
| 192.168.20.223 | 192.168.10.225 | 49681 | 49864 | TCP | 10.0m | 10.0m |
| 192.168.20.223 | 192.168.10.225 | 49688 | 49864 | TCP | 9.9m | 9.9m |
| 192.168.20.223 | 192.168.10.225 | 49696 | 49864 | TCP | 9.9m | 9.9m |
| 192.168.20.223 | 192.168.10.225 | 49692 | 49864 | TCP | 9.9m | 9.9m |
| 192.168.20.223 | 192.168.10.225 | 49677 | 49864 | TCP | 9.8m | 9.8m |
| 192.168.20.223 | 192.168.10.225 | 49669 | 49374 | TCP | 9.7m | 9.7m |

```

Sport=Source-Port Dport=Destination-Port Pro=Protocol

```

To display the specified number of top throughput connections:

```

switch:admin> portshow lan-stats --per-flow -top 5

```

| Src-Address | Dst-Address | Sport | Dport | Pro | Tx(B/s) | Rx(B/s) |
|----------------|----------------|-------|-------|-----|---------|---------|
| 192.168.20.223 | 192.168.10.225 | 49695 | 49864 | TCP | 15.7m | 15.7m |
| 192.168.20.223 | 192.168.10.225 | 49678 | 49864 | TCP | 15.7m | 15.7m |
| 192.168.20.223 | 192.168.10.225 | 49672 | 49864 | TCP | 15.7m | 15.7m |
| 192.168.20.223 | 192.168.10.225 | 49690 | 49864 | TCP | 15.6m | 15.6m |
| 192.168.20.223 | 192.168.10.225 | 49689 | 49864 | TCP | 15.6m | 15.6m |

```

Sport=Source-Port Dport=Destination-Port Pro=Protocol

```

To display the TCP error counter values:

```
switch:admin> portshow lan-stats --per-flow -tcp
***Displaying Top 1 connections by throughput***
```

| Src-Address | Dst-Address | Sport | Dport | Pro | Tx(B/s) | Rx(B/s) |
|-----------------|----------------|-------|-------|-----|---------|---------|
| TCP TxPkt RxPkt | TxDrp RxDrp | ReTx | DpAck | OOO | RTT | FlwCtrl |
| 192.168.20.223 | 192.168.10.225 | 49679 | 49864 | TCP | 17.4m | 17.4m |
| 0 1.2m | 0 0 | 0 | 0 | 42 | | 0 |

```
-----
Sport=Source-Port Dport=Destination-Port Pro=Protocol
TxPkt=Tx-Packets Rxpkt=Rx-Packets
TxDrp=TX-Drops RxDrp=RX-Drops
ReTx=ReTransmission OOO=out-of-order
DpAck=Duplicate-Acks RTT=Round-Trip-Time(milliseconds)
FlwCtrl=Number of Flow-controls
```

To display the specified number of bottom compression connections:

```
switch:admin> portshow lan-stats --per-flow -compression -bottom 5
```

| Src-Address | Dst-Address | Sport | Dport | Pro | CTx(B) | CRx(B) | CR |
|----------------|----------------|-------|-------|-----|--------|--------|--------|
| 192.168.10.225 | 192.168.20.223 | 51300 | 53086 | TCP | 104 | 8 | 0.00:1 |
| 192.168.20.223 | 192.168.10.225 | 53087 | 49374 | TCP | 239 | 143 | 0.00:1 |
| 192.168.20.223 | 192.168.10.225 | 53094 | 51301 | TCP | 195.2m | 217.7m | 10.3:1 |
| 192.168.20.223 | 192.168.10.225 | 53099 | 51301 | TCP | 197.3m | 220.2m | 10.4:1 |
| 192.168.20.223 | 192.168.10.225 | 53097 | 51301 | TCP | 159.4m | 177.8m | 10.4:1 |

```
-----
Sport=Source-Port Dport=Destination-Port Pro=Protocol
CTx(B)=Post-Compression bytes Rx(B)=Pre-Compression bytes
CR=Compression-Ratio
```

To display aggregate throughput of well-known applications:

```
switch:admin> portshow lan-stats --per-flow -application
Aggregate Info:
```

| Port | APP | TX(B/s) | RX(B/s) |
|------|---------|---------|---------|
| 80 | HTTP | 145.5m | 120.1m |
| 8080 | HTTP | 100.2m | 112.1m |
| 69 | TFTP | 601.0m | 423.3m |
| - | Unknown | 406.5m | 406.5m |

```
-----
Sport=Source-Port Dport=Destination-Port Pro=Protocol
App=Application Name
```

To display the supported well-known applications:

```
switch:admin> portshow lan-stats --known-apps
```

| App | Port-Id(s) |
|---------------|------------------------|
| CIFS | 139,445 |
| FCIP | 3225-3226 |
| FTP | 20-21,989-990,115 |
| HTTP | 80,8080,8000-8001,3128 |
| HTTPS | 443 |
| iSCSI | 3260 |
| Isilon-SyncIQ | 5666-5667 |

```

LDAP                389,8404,636
MS-SQL              1443
MySQL               3306
NETAPP-SNAP-MIRROR 10566
NFS                 2049
ORACLE-SQL         66,1525,1521
RSYNC               873
SRDF                1748
SSH                 22
SSL-SHELL           614
TELNET              23,107,513,992
TFTP                69
VERITAS-BACKUP     6101-6102,6106,3527,1125
VTS-GRID Control   1415-1416
VTS-GRID Data      350
-----

```

To display the global LAN statistics:

```

switch:admin> portshow lan-stats --global
LAN Global stats
-----
Active TCP conn           :25
Establish TCP conn       :116
Closed TCP conn          :91
TCP Tx-Bytes             :50.3m
TCP Rx-Bytes             :24.4m
TCP Tcl-Deny conn       :0
TCP Tcl lookup fail     :0
Sync-Recv                :111
Sync-fail                :0
Drop-Bytes               :0
Drop-Pkts                :0
Stale Reset from Host   :0
Number of times
Max TCP conn exceeded as client :0
Number of times
Max TCP conn exceeded as server :0
Number of times Max TCP conn
per second exceeded as client :0
Number of times Max TCP conn
per second exceeded as server :0
Number of times
Max UDP conn exceeded    :7
Number of times
TX PDU preserve ON      :0
Number of times
RX PDU preserve ON      :0
Total IPV6 pkts         :1012
FlowControl on          :0
FlowControl off         :0
Active UDP conn         :0
Establish UDP conn      :619
Closed UDP conn         :619
UDP route lookup fail   :586
UDP PDU drops due to
PKO flow control        :12800
TX UDP PDUs             :22919200
TX UPD PDU drops        :13386
RX UDP PDUs             :1991

```

portShow

```
RX UDP Tc1 lookup fail PDUs      :85
RX UDP Tc1-Deny PDUs             :1012
Total RX UDP PDU drops           :1104
RX UDP PDU drops due to
stream flow control              :0
TX ICMP PDUs                     :0
TX ICMP PDU drops                :0
RX ICMP PDUs                     :3
RX ICMP Tc1 lookup fail PDUs     :3
RX ICMP Tc1-Deny PDUs           :0
Total RX ICMP PDU drops          :3
RX ICMP PDU drops due to
stream flow control              :0
TX ASIS IP PDUs                  :0
TX ASIS IP PDU drops             :0
RX ASIS IP PDUs                  :36
RX ASIS IP Tc1 lookup fail PDUs  :20
RX ASIS IP Tc1-Deny PDUs        :8
Total RX ASIS IP PDU drops       :28
RX ASIS IP PDU drops due to
stream flow control              :0
RX Error IP Checksum             :0
RX Error TCP Checksum            :0
RX Error MAC                     :0
RX Error CRC                     :0
RX Error Parity                  :0
RX Error Length                  :0
TX UDP pkts < 64 bytes          :0
TX UDP pkts < 128 bytes         :0
TX UDP pkts < 256 bytes         :0
TX UDP pkts < 512 bytes         :0
TX UDP pkts < 1024 bytes        :22919200
TX UDP pkts < 1500 bytes        :0
TX UDP pkts < 3000 bytes        :0
TX UDP pkts < 4500 bytes        :0
TX UDP pkts < 6000 bytes        :0
TX UDP pkts < 9000 bytes        :0
RX UDP pkts < 64 bytes          :79
RX UDP pkts < 128 bytes         :992
RX UDP pkts < 256 bytes         :80
RX UDP pkts < 512 bytes         :0
RX UDP pkts < 1024 bytes        :840
RX UDP pkts < 1500 bytes        :0
RX UDP pkts < 3000 bytes        :0
RX UDP pkts < 4500 bytes        :0
RX UDP pkts < 6000 bytes        :0
RX UDP pkts < 9000 bytes        :0
-----
```

To display detailed LAN connection statistics:

```
switch:admin> portshow lan-stats --hist-stats -detail \
  -dp dp0 -index 17
Warning: It is recommended to freeze the table when using detailed stats.

DP0 Connection Detail:(Thawed)
-----
Slot/DP | Connection:  DP0 | 5
Src IP Address: 192.78.10.242
Src Port: 53795
```



```

Dst IP Address:      192.79.10.243
Dst Port:           59039
Start time:         08/18/15 21:17:26 UTC
End time:           08/18/15 21:17:59 UTC
Close reason/Flag:  Remote LAN / TX FIN
Current/Previous State: Time-Wait / FIN-Wait-2
Last 5 Connection States:
  State/Reason Code 1: SYN-Sent / User Connect
  State/Reason Code 2: Established / Connection Success
  State/Reason Code 3: FIN-Wait-1 / User Close
  State/Reason Code 4: FIN-Wait-2 / RX ACK for FIN
  State/Reason Code 5: Time-Wait / Timer Cleanup
L4 Protocol/MSS:   TCP / 1460
TX Stats
  TX Bytes/Packets: 710717432 / 542617
  Slow Starts:      0
  FastRetrans/RetransTO: 0 / 0
  Initial/Max Send Seq: 1490083657 / 2179096399
RX Stats
  RX Bytes/Packets: 689012740 / 504638
  Out of Orders/Dup Ack: 0 / 0
  Initial Receive Seq: 1799596954

```

To freeze the LAN statistic table:

portshow lan-stats --hist-stats -freeze

DPO Connection Summary:(Frozen)

| Idx | Src-Address | Dst-Address | Sport | Dport | Pro | Tx(B) | Rx(B) |
|-------------------------|-------------------|---------------|-------|-------|-----|--------|--------|
| ----- | | | | | | | |
| First 5 Connections: | | | | | | | |
| 0 | 192.78.10.242 | 192.79.10.243 | 53786 | 49374 | TCP | 271 | 72 |
| 1 | 192.78.10.242 | 192.79.10.243 | 53785 | 59038 | TCP | 120 | 8 |
| 2 | 192.78.10.242 | 192.79.10.243 | 53788 | 59039 | TCP | 711.1m | 689.4m |
| 3 | 192.78.10.242 | 192.79.10.243 | 53787 | 59039 | TCP | 712.3m | 690.5m |
| 4 | 192.78.10.242 | 192.79.10.243 | 53791 | 59039 | TCP | 713.0m | 691.2m |
| Last 5 Connections: | | | | | | | |
| 27 | 192.78.10.242 | 192.79.10.243 | 53816 | 59039 | TCP | 697.2m | 675.9m |
| 28 | 192.78.10.242 | 192.79.10.243 | 53814 | 59039 | TCP | 701.1m | 679.7m |
| 29 | 192.78.10.242 | 192.79.10.243 | 53815 | 59039 | TCP | 700.4m | 679.0m |
| 30 | 192.78.10.242 | 192.79.10.243 | 53810 | 59039 | TCP | 707.7m | 686.1m |
| 31 | 192.78.10.242 | 192.79.10.243 | 53817 | 59039 | TCP | 691.0m | 669.9m |
| ----- | | | | | | | |
| Total Connection count: | 32 | | | | | | |
| Oldest Entry: | 08/18/15 21:17:59 | | | | | | |
| Newest Entry: | 08/18/15 21:17:59 | | | | | | |
| Close RX/TX FIN: | 1 / 31 | | | | | | |
| Close RX/TX RST: | 0 / 0 | | | | | | |
| Total TX Errors | | | | | | | |
| Slow Starts: | 0 | | | | | | |
| FastRetrans/RetransTO: | 0 / 0 | | | | | | |
| Total RX Errors | | | | | | | |
| Out of Orders/Dup Ack: | 0 / 0 | | | | | | |

To thaw the LAN statistic table:

portshow lan-stats --hist-stats -thaw

DP0 Connection Summary:(Thawed)

```

-----
  Idx  Src-Address          Dst-Address          Sport  Dport  Pro  Tx(B)    Rx(B)
-----
First 5 Connections:
  0    192.78.10.242         192.79.10.243       53786  49374  TCP  271      72
  1    192.78.10.242         192.79.10.243       53785  59038  TCP  120      8
  2    192.78.10.242         192.79.10.243       53788  59039  TCP  711.1m   689.4m
  3    192.78.10.242         192.79.10.243       53787  59039  TCP  712.3m   690.5m
  4    192.78.10.242         192.79.10.243       53791  59039  TCP  713.0m   691.2m
Last 5 Connections:
  27   192.78.10.242         192.79.10.243       53816  59039  TCP  697.2m   675.9m
  28   192.78.10.242         192.79.10.243       53814  59039  TCP  701.1m   679.7m
  29   192.78.10.242         192.79.10.243       53815  59039  TCP  700.4m   679.0m
  30   192.78.10.242         192.79.10.243       53810  59039  TCP  707.7m   686.1m
  31   192.78.10.242         192.79.10.243       53817  59039  TCP  691.0m   669.9m
-----
Total Connection count:    32
Oldest Entry:              08/18/15 21:17:58
Newest Entry:              08/18/15 21:17:58
Close RX/TX FIN:          1 / 31
Close RX/TX RST:          0 / 0
Total TX Errors
  Slow Starts:              0
  FastRetrans/RetransTO:   0 / 0
Total RX Errors
  Out of Orders/Dup Ack:   0 / 0

```

To display the SLA statistics:

```

switch:admin> portshow sla
  Name  PktLoss  Runtime  Timeout
-----
  netA  0.50%    15m      1h
-----

```

To display all tunnels and only circuits using IP address 10.0.0.1:

```

switch:admin> portshow fcipunnel -c
                --filter -ipaddr 10.0.0.1 --default show

```

To display only per-flow statistics that use both IP address 192.168.0.10 AND TCP port 336:

```

switch:admin> portshow lan-stats --per-flow -all
                --filter -ipaddr 192.168.0.10 -tcp-port 336 -and

```

Function Display FCIP tunnels, circuits, management interfaces and FICON statistics on the Brocade 7840, SX6, or FX8-24 platforms:

Synopsis `portshow option [all | [slot/]ve_port] arguments [optional_arguments]`

Description Use this command to display configuration parameters and status information for FCIP tunnels and FCIP circuits on the GbE/10GbE ports on the Brocade 7840, Brocade SX6, and FX8-24 platforms. The FCIP tunnels on the local and remote GbE ports act as Virtual E_Ports (VE_Ports) connecting the local and remote fabrics. The following display options are supported on these platforms:

- **portShow fciptunnel** - Displays Fibre Channel over IP (FCIP) tunnels including the following:
 - Additional performance information
 - FICON configuration parameters
 - IPSec status (enabled or disabled), mode (legacy if configured) and key
 - Circuits with the tunnel
 - QoS statistics including performance for each priority
 - FCIP Tunnel hierarchy listing circuit IDs and TCP connection information for the tunnel
- **portShow fcipcircuit** - Displays status and configuration for FCIP circuits:
 - TCP statistics for the circuit
 - Circuit-level VLAN tagging configuration
 - Circuit-level Differentiated Services (DSCP) markings
- **portShow xtun** - Displays FICON and FCP emulation statistics and current runtime conditions.

Operands This command has the following operands:

| | |
|-----------------------|--|
| <i>slot</i> | For bladed systems only, specifies the slot number of the VE_Port to be displayed, followed by a slash (/). |
| <i>ve_port</i> | Displays information for a single specified FCIP tunnel. On the Brocade 7840, SX6, or FX8-24, specify the VE_Port number associated with the tunnel configured on one of the GbE ports. VE_Ports are numbered 12-31 on the Brocade FX8-24 blade. |
| all | Displays information for all configured FCIP tunnels. |
| fciptunnel | Displays configuration and status per FCIP tunnel. The following operands are supported with fciptunnel and fcipcircuit . |
| -c --circuit | Displays the FCIP circuits within the tunnel. This operand can be used with any other operand to include circuit displays. |
| -t --tcp | Displays the TCP statistics for the circuit. The -c option must be specified with this option. |
| --reset | Displays time based set of statistics. |
| --lifetime | Displays the entire lifetime statistics for FCIP Tunnels, Circuits and the associated TCP connections. This option will ignore any time based deltas that were created previously set using the --reset option. |

| | |
|--------------------------|---|
| -p --perf | Displays additional performance information for the specified FCIP tunnels. |
| --lifetime | Displays the entire lifetime statistics for FCIP Tunnels, Circuits and the associated TCP connections. This option will ignore any time based deltas that were created previously set using the --reset option. |
| -q --qos | Displays the QoS statistics including performance for each priority. This operand is optional with the --perf option. |
| -C --config | Displays only the configuration information |
| -i --ipsec | Displays IPsec status (enabled or disabled) and key if enabled. If IPsec is enabled and configured in legacy mode, the mode information is displayed in parenthesis. This option is valid for portShow fcip tunnel only. This option is deprecated but still recognized. Use the -C --config option instead. |
| -h --hier | Displays the FCIP tunnel hierarchy listing basic circuit IDs and TCP connection information for the tunnel. |
| -s --summary | Displays a summary view of the tunnel configuration parameters for a specific VE_Port. You can use the summary option with the -perf and the -qos option. When used with the -perf option, the summary option forces the --circuit option. Flags indicate ipsec-configured tunnels, legacy ipsec tunnels, and compression mode. Circuit flags indicate vlan tagging, crossport configuration, and ipv4 or ipv6 configuration. |
| -d --detail | Displays a full view configuration details for all configured tunnels when used with the all port specifier. This view provides an alternative to specifying a specific VE_Port or to using the --perf option for a comprehensive view. |
| -l --ip-address | Displays the IP addresses configured the specified circuits. You must use this option with either the --summary or the --circuit option. This option is deprecated but still recognized. Use the -C --config option instead. |
| fcip circuit | Displays FCIP circuit configuration and status. The following operands are supported with portShow fcip circuit : |
| <i>circuit_ID</i> | Specifies a single circuit within the FCIP tunnel. The circuit ID is an integer value between 1 and 19 on the Brocade FX8-24. This operand is required. |
| -t --tcp | Displays the TCP statistics for the specified circuit. |
| --reset | Displays time based set of statistics. |

| | |
|----------------------|---|
| --lifetime | Displays the entire lifetime statistics for FCIP Tunnels, Circuits and the associated TCP connections. This option will ignore any time based deltas that were created previously set using the --reset option. |
| -p --perf | Displays additional performance information for the specified FCIP circuit. |
| --lifetime | Displays the entire lifetime statistics for FCIP Tunnels, Circuits and the associated TCP connections. This option will ignore any time based deltas that were created previously set using the --reset option. |
| -q --qos | Displays the QoS statistics including performance for each priority. This operand is optional with the --perf option. |
| -C --config | Displays only the configuration information |
| xtun | Displays FICON and FCP emulation statistics and current runtime conditions for a specified set of parameters. The following arguments are supported: |
| -fcp | Displays the SCSI FastWrite/Tape Pipelining command sub-menu when issued with a VE_Port number. The syntax for -fcp is as follows: portshow xtun [slot/] \ve_port -fcp [level] [command] [param] [options] The following optional operands are supported with -fcp to display FCP emulation statistics and status information: |
| -help | Displays the command usage. You must specify a VE_Port number to display the help functions, for example: portshow xtun 7/12 -fcp -help . |
| level | Specifies the level for which information is displayed. You can specify one or more of the following levels. With each additional level, this command generates progressively more information. |
| -port | Displays data at the port level. |
| -it | Displays data at the Initiator Target (SID/DID) FCP level. |
| -itn | Displays data at the Initiator Target nexus (SID/DID) level. |
| -itl | Displays data at the Initiator Target LUN (SID/DID/LUN) level. |
| -twb | Displays data at the Exchange (SID/DID/LUN/Exchange) level. |

command

Specifies the type of information to be displayed. This operand is optional; if omitted, the default (**-stats**) is used. You can specify more than one command option. Valid commands include the following:

-stats

Displays FCP emulation statistics. This is the default display.

-info

Displays general FCP emulation information.

-cfg

Displays the FCP emulation configuration.

-dump

Displays a raw data dump including data for all information types.

param

Limits output to one or more of the following parameters, given the commands and levels specified with this command. This operand is optional; if omitted, output for all parameters is displayed. There is no default parameter.

-sid *SID*

Displays output for the specified SID only.

-did *DID*

Displays output for the specified DID only.

-lun *LUN*

Displays output for the specified LUN only.

-timer

Displays timer information only.

options

Affects all levels and commands globally. Only one option is supported.

-zero

Displays zero-valued statistics. Note that some commands may show zero-valued information regardless of whether or not this option is specified.

-ficon

Displays sub-menu for FICON emulation display commands when issued with a VE_Port number. The syntax for **-ficon** is as follows:

```
portshow xtun [slot/]ve_port -ficon [command] [options]
```

The following optional commands are supported with **-ficon**; if omitted, the usage for all parameters is displayed. The **xtun -ficon** command options include displays for all types of FICON Emulation.

-help

Displays the command usage. You must specify a VE_Port number to display the help functions, for example, **portshow xtun 7/12 -ficon -help**.

| | |
|--------------------------|---|
| -stats | Displays global FICON Emulation statistics for the tunnel including FICON XRC Emulation, FICON Tape Write, FICON Tape Read, FICON Teradata Write, and FICON Teradata Read statistics. |
| -fdpb <i>adrs</i> | Displays FICON ports or a specific FICON Device Path Block. |
| -fchb <i>adrs</i> | Displays FICON logical partitions (LPARs) or a specific FICON Channel Control Block. |
| -fcub <i>adrs</i> | Displays FICON images (the same output as with -images) or a specific FICON Control Unit Block. |
| -images | Displays FICON images. |
| -fdcb <i>adrs</i> | Displays FICON devices or specific FICON Device Control Block. |
| -tapeperf | Starts the emulated Tape Read and Write performance monitor or displays the performance statistics. when you first issue this command or any other performance monitor commands after a reboot, the command starts the performance monitor, takes a snapshot of current statistics, and saves them with a time stamp. When you issue the command again, it displays the time elapsed between the two iterations of the command and the average time delta statistics. A new time stamp and current statistics are saved as a basis for the next iteration. |
| -teraperf | Starts the emulated Teradata performance monitor or displays the performance statistics. |
| -printperf | Starts the emulated Printer performance monitor or displays the performance statistics. |
| -xrcperf | Starts the emulated XRC performance monitor or displays the performance statistics. |
| -structs | Displays FICON control block sizes. |
| -emul | Displays comprehensive FICON emulation statistics. Use one of the following options to display emulation statistics about a specific component. |
| -emulxrc | Displays FICON XRC emulation statistics. |
| -emultape | Displays FICON Tape emulation statistics. |

portShow

- emultera** Displays FICON Teradata emulation statistics.
- emulprint** Displays FICON Printer emulation statistics.
- act** Displays the current Active Exchange information.
- options** The following additional option is supported.
 - clear** Resets the specified statistics. This operand is optional; it requires a preceding command.
- mem *adrs length*** Specifies the SE memory length in words to display. The valid range for *length* is 0 to 1024.
- dram2** Display the current usage of the dynamic memory allocator.
- pools** Display the current free pool allocation (FPA) buffer usage.
- tcb** Displays tunnel statistics.
- drshow** Displays the current Descriptor Ring status
- smem** Displays the current Shadow Memory data.
- rte** Displays the current Routing Info SE memory.
- filter *filter_args*** Filters the **portShow** output based on the specified filter arguments. The filter arguments can be specified in any combination or as a conditional statement using the logical AND or OR operator. A conditional statement can have up to 30 conditions. The following filter arguments are supported:
 - ipaddr *ip_address*** Filters the output based on the specified IP address.
 - port [*slot*]*port*** Filters the output based on the specified port number.
 - tcp-port *port*** Filters the output based on the specified TCP port number.
 - dp [*slot*]*dp#*** Filters the output based on the specified dual processor ID.
 - retransmits *value*** Filters the output based on the retransmits exceeding specified value.

- rtt *ms***
Filters the output based on the specified circuit round trip time in milliseconds.
- connect-count *value***
Filters the output based on the connected count exceeding specified value.
- bps *value*[*k* | *m*]**
Filters the output based on bandwidth (bytes per second) exceeding the specified value.
Specify **k** for KBps and **m** for MBps.
- or**
The logical OR operator.
- and**
The logical AND operator.
- default show | hide**
Sets the default display action if the specified filter statement is not supported. The default action is **hide**.

Examples To display FCIP tunnel configuration parameters on the Brocade SX6 blade with Ficon enabled:

```
switch:admin> portshow fciptunnel 8/27 -d

Tunnel: VE-Port:8/27 (idx:11, DP1)
=====
Oper State           : Online
TID                  : 219
Flags                : 0x00000000
IP-Extension         : Disabled
Compression          : None
QoS BW Ratio         : 50% / 30% / 20%
Fastwrite            : Enabled
Tape Pipelining      : Read/Write
IPSec                : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN            : 10:00:00:05:33:e7:d2:13
Peer WWN             : 10:00:00:05:33:65:a8:05
RemWWN (config)     : 00:00:00:00:00:00:00:00
FICON                : Enabled
  XRC                 : Enabled
  Tape Pipelining     : Read/Write
  Read Block id       : Enabled
  Tin Tir             : Enabled
  DevAck              : Enabled
  Tera Read           : Disabled
  Tera Write          : Disabled
  Print               : Disabled
FICON Max Read Pipe  : 32
FICON Max Write Pipe : 32
FICON Max Read Devs  : 16
FICON Max Write Devs : 16
FICON Write Timer    : 300
FICON Write Chain    : 3200000
FICON OXID Base      : 0x8000
FICON Debug Flags    : 0x77c90000
cfgmask              : 0x003fc7ff 0x40000208
Flow Status          : 0
ConCount/Duration    : 1 / 35m19s
```

portShow

```
Uptime           : 50m13s
Stats Duration   : 35m19s
Receiver Stats   : 846212584 bytes / 4594438 pkts / 560.99 KBps Avg
Sender Stats     : 216803637504 bytes / 29122536 pkts / 144.36 MBps Avg
TCP Bytes In/Out : 3135072280 / 223666014748
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max) : 1 / 1 / 1 ms
Wan Util         : 26.3%
TxQ Util         : 0.0%
```

To display an FCIP tunnel with FICON disabled:

```
switch:admin> portshow fciptunnel 8/26 -d
```

```
Tunnel: VE-Port:8/26 (idx:10, DP1)
=====
Oper State       : Disabled
TID              : 218
Flags            : 0x00000000
IP-Extension     : Disabled
Compression      : None
QoS BW Ratio     : 50% / 30% / 20%
Fastwrite        : Disabled
Tape Pipelining  : Disabled
IPSec            : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN        : 10:00:00:05:33:e7:d2:11
Peer WWN         : 00:00:00:00:00:00:00:00
RemWWN (config) : 00:00:00:00:00:00:00:00
cfgmask         : 0x000004ff 0x40000208
Flow Status      : 0
ConCount/Duration : 0 / 1h6m7s
Uptime          : 0s
Stats Duration   : 0s
Receiver Stats   : 0 bytes / 0 pkts / 0.00 Bps Avg
Sender Stats     : 0 bytes / 0 pkts / 0.00 Bps Avg
TCP Bytes In/Out : 0 / 0
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max) : 0 / 0 / 0 ms
Wan Util         : 0.0%
TxQ Util         : 0.0%
```

To display an FCIP tunnel with additional circuit information:

```
switch:admin> portshow fciptunnel 8/27 -d
```

```
Tunnel: VE-Port:8/27 (idx:11, DP1)
=====
Oper State       : Online
TID              : 219
Flags            : 0x00000000
IP-Extension     : Disabled
Compression      : None
QoS BW Ratio     : 50% / 30% / 20%
Fastwrite        : Enabled
Tape Pipelining  : Read/Write
IPSec            : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN        : 10:00:00:05:33:e7:d2:13
```

```

Peer WWN           : 10:00:00:05:33:65:a8:05
RemWWN (config)   : 00:00:00:00:00:00:00:00
FICON              : Enabled
  XRC              : Enabled
  Tape Pipelining : Read/Write
  Read Block id   : Enabled
  Tin Tir         : Enabled
  DevAck          : Enabled
  Tera Read       : Disabled
  Tera Write      : Disabled
  Print           : Disabled
FICON Max Read Pipe : 32
FICON Max Write Pipe : 32
FICON Max Read Devs : 16
FICON Max Write Devs : 16
FICON Write Timer   : 300
FICON Write Chain   : 3200000
FICON OXID Base     : 0x8000
FICON Debug Flags   : 0x77c90000
cfgmask            : 0x003fc7ff 0x40000208
Flow Status         : 0
ConCount/Duration  : 1 / 35m19s
Uptime             : 50m13s
Stats Duration     : 35m19s
Receiver Stats     : 846212584 bytes / 4594438 pkts / 560.99 KBps Avg
Sender Stats       : 216803637504 bytes / 29122536 pkts / 144.36 MBps Avg
TCP Bytes In/Out   : 3135072280 / 223666014748
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)  : 1 / 1 / 1 ms
Wan Util           : 26.3%
TxQ Util           : 0.0%

```

To display additional performance parameters on tunnel 1/12 (add -c to display all circuits):

```
switch:admin> portshow fcipunnel 8/27 -c
```

```

Tunnel: VE-Port:8/27 (idx:11, DP1)
=====
Oper State           : Online
TID                  : 219
Flags                : 0x00000000
IP-Extension         : Disabled
Compression          : None
QoS BW Ratio         : 50% / 30% / 20%
Fastwrite            : Enabled
Tape Pipelining     : Read/Write
IPSec                : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN            : 10:00:00:05:33:e7:d2:13
Peer WWN             : 10:00:00:05:33:65:a8:05
RemWWN (config)     : 00:00:00:00:00:00:00:00
FICON                : Enabled
  XRC                : Enabled
  Tape Pipelining   : Read/Write
  Read Block id    : Enabled
  Tin Tir          : Enabled
  DevAck           : Enabled
  Tera Read        : Disabled
  Tera Write       : Disabled
  Print            : Disabled

```

```

FICON Max Read Pipe : 32
FICON Max Write Pipe : 32
FICON Max Read Devs : 16
FICON Max Write Devs : 16
FICON Write Timer : 300
FICON Write Chain : 3200000
FICON OXID Base : 0x8000
FICON Debug Flags : 0x77c90000
cfgmask : 0x003fc7ff 0x40000208
Flow Status : 0
ConCount/Duration : 1 / 33m59s
Uptime : 48m53s
Stats Duration : 33m59s
Receiver Stats : 811663752 bytes / 4406845 pkts / 554.94 KBps Avg
Sender Stats : 207944106140 bytes / 27935226 pkts / 143.31 MBps Avg
TCP Bytes In/Out : 3005648720 / 214490692956
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max) : 1 / 1 / 1 ms
Wan Util : 24.3%
TxQ Util : 0.0%

```

Circuit 8/27.0 (DP1)

```

=====
Admin/Oper State : Enabled / Online
Flags : 0x00000000
IP Addr (L/R) : 192.168.2.170 8/ge2 <-> 192.168.2.164
HA IP Addr (L/R) : 192.168.2.70 8/ge2 <-> 192.168.2.64
Configured Comm Rates: 2000000 / 5000000 kbps
Peer Comm Rates : 2000000 / 5000000 kbps
Actual Comm Rates : 2000000 / 5000000 kbps
Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
Metric : 0
Connection Type : Default
ARL-Type : Reset
PMTU : Disabled
SLA : (none)
Failover Group : 0
VLAN-ID : NONE
L2Cos (FC:h/m/l) : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l) : 0 / 0 / 0
DSCP (FC:h/m/l) : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l) : 0 / 0 / 0
cfgmask : 0x40000000 0x00013c2f
Flow Status : 0
ConCount/Duration : 1 / 33m59s
Uptime : 48m53s
Stats Duration : 33m59s
Receiver Stats : 405802244 bytes / 2203254 pkts / 277.42 KBps Avg
Sender Stats : 103969139456 bytes / 13967446 pkts / 71.69 MBps Avg
TCP Bytes In/Out : 1502340852 / 107217064500
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max) : 1 / 1 / 1 ms
Wan Util : 24.5%

```

Circuit 8/27.1 (DP1)

```

=====
Admin/Oper State : Enabled / Online
Flags : 0x00000000
IP Addr (L/R) : 192.168.3.170 8/ge3 <-> 192.168.3.164
HA IP Addr (L/R) : 192.168.3.70 8/ge3 <-> 192.168.3.64

```

```

Configured Comm Rates: 2000000 / 5000000 kbps
Peer Comm Rates       : 2000000 / 5000000 kbps
Actual Comm Rates     : 2000000 / 5000000 kbps
Keepalive (Cfg/Peer)  : 1000 (1000 / 1000) ms
Metric                : 0
Connection Type       : Default
ARL-Type              : Reset
PMTU                  : Disabled
SLA                   : (none)
Failover Group        : 0
VLAN-ID               : NONE
L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)     : 0 / 0 / 0
DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l)      : 0 / 0 / 0
cfgmask               : 0x40000000 0x00013c2f
Flow Status           : 0
ConCount/Duration     : 1 / 33m59s
Uptime                : 48m53s
Stats Duration        : 33m59s
Receiver Stats        : 405861508 bytes / 2203591 pkts / 277.52 KBps Avg
Sender Stats          : 103974966684 bytes / 13967780 pkts / 71.62 MBps Avg
TCP Bytes In/Out      : 1503305492 / 107273625304
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util              : 24.1%

```

To display TCP connections for the circuits on the tunnel:

```
switch:admin> portshow fciptunnel 8/27 -circuit
```

```

Tunnel: VE-Port:8/27 (idx:11, DP1)
=====
Oper State           : Online
TID                  : 219
Flags                : 0x00000000
IP-Extension         : Disabled
Compression          : None
QoS BW Ratio         : 50% / 30% / 20%
Fastwrite            : Enabled
Tape Pipelining      : Read/Write
IPSec                : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN             : 10:00:00:05:33:e7:d2:13
Peer WWN              : 10:00:00:05:33:65:a8:05
RemWWN (config)      : 00:00:00:00:00:00:00:00
FICON                : Enabled
  XRC                 : Enabled
  Tape Pipelining     : Read/Write
  Read Block id       : Enabled
  Tin Tir             : Enabled
  DevAck              : Enabled
  Tera Read           : Disabled
  Tera Write          : Disabled
  Print               : Disabled
FICON Max Read Pipe  : 32
FICON Max Write Pipe : 32
FICON Max Read Devs  : 16
FICON Max Write Devs : 16
FICON Write Timer    : 300

```

```

FICON Write Chain      : 3200000
FICON OXID Base       : 0x8000
FICON Debug Flags     : 0x77c90000
cfgmask               : 0x003fc7ff 0x40000208
Flow Status           : 0
ConCount/Duration     : 1 / 36m53s
Uptime                : 51m46s
Stats Duration        : 36m52s
Receiver Stats        : 888449704 bytes / 4823789 pkts / 518.43 KBps Avg
Sender Stats          : 227615878740 bytes / 30565366 pkts / 133.07 MBps Avg
TCP Bytes In/Out      : 3293193644 / 234818572912
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util              : 31.1%
TxQ Util              : 0.0%

```

Circuit 8/27.0 (DP1)

```

=====
Admin/Oper State      : Enabled / Online
Flags                 : 0x00000000
IP Addr (L/R)         : 192.168.2.170 8/ge2 <-> 192.168.2.164
HA IP Addr (L/R)     : 192.168.2.70 8/ge2 <-> 192.168.2.64
Configured Comm Rates: 2000000 / 5000000 kbps
Peer Comm Rates       : 2000000 / 5000000 kbps
Actual Comm Rates     : 2000000 / 5000000 kbps
Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
Metric               : 0
Connection Type       : Default
ARL-Type              : Reset
PMTU                  : Disabled
SLA                   : (none)
Failover Group        : 0
VLAN-ID               : NONE
L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)     : 0 / 0 / 0
DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l)      : 0 / 0 / 0
cfgmask               : 0x40000000 0x00013c2f
Flow Status           : 0
ConCount/Duration     : 1 / 36m53s
Uptime                : 51m47s
Stats Duration        : 36m53s
Receiver Stats        : 444186300 bytes / 2411679 pkts / 259.59 KBps Avg
Sender Stats          : 113803731376 bytes / 15282510 pkts / 66.53 MBps Avg
TCP Bytes In/Out      : 1646503476 / 117412356592
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util              : 31.4%

```

TCP Connection 8/27.0 HA-Type:Local-Backup Pri:Low Conn:0x055665f4

```

=====
Local / Remote Port   : 3225 / 55379
Duration              : 36m52s
MSS                   : 1460 bytes
ARL Min / Cur / Max   : 33328 / 33328 / 833328
ARL Reset Algo        : Reset
Send Window
Size / Scale          : 20832256 / 9
Slow Start Threshold : 16777216
Congestion Window     : 16778676

```

```

    Pkts InFlight          : 0
  Recv Window
    Size / Scale           : 20832768 (Max:20832768) / 9
  SendQ Next / Min / Max  : 0x8e5c8b31 / 0x8e5c8b31 / 0x8e5c8b31
  RecvQ Next / Min / Max  : 0x0e776252 / 0x0e776252 / 0x0fb54402
  RecvQ Pkts              : 0
  Sender Stats
    Sent Bytes / Pkts      : 510608 / 5742
    Unacked Data          : 0
    Retransmits Slow / Fast : 0 / 0 (High:0)
    SlowStart              : 0
  Reciever Stats
    Recv Bytes / Pkts      : 536792 / 5742
    Out-of-Order           : 0 (High:0)
    Duplicate ACKs         : 0
    RTT / Variance (High)  : 0 ms (0 ms) / 0 ms (0 ms)
[output truncated]

```

To display the time based set of statistics:

```
switch:admin> portshow fciptunnel 8/27 -cp
```

```

Tunnel: VE-Port:8/27 (idx:11, DP1)
=====
Oper State           : Online
TID                  : 219
Flags                : 0x00000000
IP-Extension         : Disabled
Compression          : None
QoS BW Ratio         : 50% / 30% / 20%
Fastwrite            : Enabled
Tape Pipelining      : Read/Write
IPSec                : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN             : 10:00:00:05:33:e7:d2:13
Peer WWN              : 10:00:00:05:33:65:a8:05
RemWWN (config)      : 00:00:00:00:00:00:00:00
FICON                : Enabled
  XRC                 : Enabled
  Tape Pipelining     : Read/Write
  Read Block id       : Enabled
  Tin Tir             : Enabled
  DevAck              : Enabled
  Tera Read           : Disabled
  Tera Write          : Disabled
  Print               : Disabled
FICON Max Read Pipe  : 32
FICON Max Write Pipe : 32
FICON Max Read Devs  : 16
FICON Max Write Devs : 16
FICON Write Timer    : 300
FICON Write Chain    : 3200000
FICON OXID Base      : 0x8000
FICON Debug Flags    : 0x77c90000
cfgmask              : 0x003fc7ff 0x40000208
Flow Status          : 0
ConCount/Duration    : 1 / 38m38s
Uptime               : 53m32s
Stats Duration       : 38m38s
Receiver Stats       : 4495871308 bytes / 5361838 pkts / 88.46 MBps Avg

```

portShow

```
Sender Stats          : 230389776820 bytes / 31031771 pkts / 14.02 MBps Avg
TCP Bytes In/Out     : 7012468568 / 237786788876
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util              : 3.2%
TxQ Util              : 0.0%
```

Circuit 8/27.0 (DP1)

```
=====
Admin/Oper State     : Enabled / Online
Flags                 : 0x00000000
IP Addr (L/R)        : 192.168.2.170 8/ge2 <-> 192.168.2.164
HA IP Addr (L/R)     : 192.168.2.70 8/ge2 <-> 192.168.2.64
Configured Comm Rates: 2000000 / 5000000 kbps
Peer Comm Rates      : 2000000 / 5000000 kbps
Actual Comm Rates    : 2000000 / 5000000 kbps
Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
Metric                : 0
Connection Type      : Default
ARL-Type             : Reset
PMTU                  : Disabled
SLA                   : (none)
Failover Group       : 0
VLAN-ID              : NONE
L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)     : 0 / 0 / 0
DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l)      : 0 / 0 / 0
cfgmask               : 0x40000000 0x00013c2f
Flow Status          : 0
ConCount/Duration    : 1 / 38m38s
Uptime                : 53m33s
Stats Duration       : 38m38s
Receiver Stats       : 2250048596 bytes / 2680666 pkts / 44.28 MBps Avg
Sender Stats         : 115191206632 bytes / 15515761 pkts / 7.02 MBps Avg
TCP Bytes In/Out     : 3497989220 / 118885996076
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util              : 3.2%
```

Circuit 8/27.1 (DP1)

```
=====
Admin/Oper State     : Enabled / Online
Flags                 : 0x00000000
IP Addr (L/R)        : 192.168.3.170 8/ge3 <-> 192.168.3.164
HA IP Addr (L/R)     : 192.168.3.70 8/ge3 <-> 192.168.3.64
Configured Comm Rates: 2000000 / 5000000 kbps
Peer Comm Rates      : 2000000 / 5000000 kbps
Actual Comm Rates    : 2000000 / 5000000 kbps
Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
Metric                : 0
Connection Type      : Default
ARL-Type             : Reset
PMTU                  : Disabled
SLA                   : (none)
Failover Group       : 0
VLAN-ID              : NONE
L2Cos (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)     : 0 / 0 / 0
DSCP (FC:h/m/l)      : 0 / 0 / 0 (Ctrl:0)
```



```

DSCP (IP:h/m/l)      : 0 / 0 / 0
cfgmask              : 0x40000000 0x00013c2f
Flow Status          : 0
ConCount/Duration    : 1 / 38m38s
Uptime               : 53m33s
Stats Duration       : 38m38s
Receiver Stats       : 2245929472 bytes / 2681187 pkts / 44.17 MBps Avg
Sender Stats         : 115198570188 bytes / 15516010 pkts / 7.00 MBps Avg
TCP Bytes In/Out     : 3514474436 / 118900788176
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 1 / 1 / 1 ms
Wan Util             : 3.2%

```

To display the entire lifetime statistics for FCIP Tunnels, Circuits and the associated TCP connections:

```
switch:admin> portshow fciptunnel 8/27 -ctL
```

```

Tunnel: VE-Port:8/27 (idx:11, DP1)
=====
Oper State           : Online
TID                  : 219
Flags                 : 0x00000000
IP-Extension         : Disabled
Compression          : None
QoS BW Ratio         : 50% / 30% / 20%
Fastwrite            : Enabled
Tape Pipelining      : Read/Write
IPSec                : Disabled
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN            : 10:00:00:05:33:e7:d2:13
Peer WWN             : 10:00:00:05:33:65:a8:05
RemWWN (config)     : 00:00:00:00:00:00:00:00
FICON                : Enabled
  XRC                 : Enabled
  Tape Pipelining     : Read/Write
  Read Block id       : Enabled
  Tin Tir             : Enabled
  DevAck              : Enabled
  Tera Read           : Disabled
  Tera Write          : Disabled
  Print               : Disabled
FICON Max Read Pipe  : 32
FICON Max Write Pipe : 32
FICON Max Read Devs  : 16
FICON Max Write Devs : 16
FICON Write Timer    : 300
FICON Write Chain    : 3200000
FICON OXID Base      : 0x8000
FICON Debug Flags    : 0x77c90000
cfgmask              : 0x003fc7ff 0x40000208
Flow Status          : 0
ConCount/Duration    : 2 / 1h15m52s
Uptime               : 54m58s
Stats Duration       : 54m58s
Receiver Stats       : 12183885664 bytes / 7267213 pkts / 86.35 MBps Avg
Sender Stats         : 272470084404 bytes / 37117428 pkts / 13.91 MBps Avg
TCP Bytes In/Out     : 15336120428 / 281269480752
ReTx/OOO/SloSt/DupAck: 35 / 0 / 20 / 0
RTT (min/avg/max)    : 1 / 1 / 11 ms
Wan Util             : 3.2%

```

TxQ Util : 0.0%

Circuit 8/27.0 (DP1)

```

=====
Admin/Oper State      : Enabled / Online
Flags                 : 0x00000000
IP Addr (L/R)        : 192.168.2.170 8/ge2 <-> 192.168.2.164
HA IP Addr (L/R)     : 192.168.2.70 8/ge2 <-> 192.168.2.64
Configured Comm Rates: 2000000 / 5000000 kbps
Peer Comm Rates      : 2000000 / 5000000 kbps
Actual Comm Rates    : 2000000 / 5000000 kbps
Keepalive (Cfg/Peer) : 1000 (1000 / 1000) ms
Metric               : 0
Connection Type      : Default
ARL-Type             : Reset
PMTU                 : Disabled
SLA                  : (none)
Failover Group       : 0
VLAN-ID              : NONE
L2Cos (FC:h/m/l)    : 0 / 0 / 0 (Ctrl:0)
L2Cos (IP:h/m/l)    : 0 / 0 / 0
DSCP (FC:h/m/l)     : 0 / 0 / 0 (Ctrl:0)
DSCP (IP:h/m/l)     : 0 / 0 / 0
cfgmask              : 0x40000000 0x00013c2f
Flow Status          : 0
ConCount/Duration    : 2 / 1h15m52s
Uptime               : 54m59s
Stats Duration       : 54m59s
Receiver Stats       : 6093629652 bytes / 3633882 pkts / 43.17 MBps Avg
Sender Stats         : 136234172252 bytes / 18558931 pkts / 6.94 MBps Avg
TCP Bytes In/Out     : 9506571664 / 191565449080
ReTx/OOO/SloSt/DupAck: 33 / 0 / 18 / 0
RTT (min/avg/max)   : 1 / 1 / 11 ms
Wan Util             : 3.2%

```

TCP Connection 8/27.0 HA-Type:Local-Backup Pri:Low Conn:0x055665f4

```

=====
Local / Remote Port   : 3225 / 55379
Duration              : 46m42s
MSS                   : 1460 bytes
ARL Min / Cur / Max   : 33328 / 33328 / 833328
ARL Reset Algo        : Reset
Send Window
  Size / Scale        : 20832256 / 9
  Slow Start Threshold : 16777216
  Congestion Window   : 16778676
  Pkts InFlight       : 0
Recv Window
  Size / Scale        : 20832768 (Max:20832768) / 9
SendQ Next / Min / Max : 0x8e5cea51 / 0x8e5cea51 / 0x8e5cea51
RecvQ Next / Min / Max : 0x0e77cb0a / 0x0e77cb0a / 0x0fb5acba
RecvQ Pkts           : 0
Sender Stats
  Sent Bytes / Pkts   : 650880 / 7292
  Unacked Data        : 0
  Retransmits Slow / Fast : 0 / 0 (High:0)
  SlowStart           : 0

```

```

Receiver Stats
  Recv Bytes / Pkts      : 682504 / 7292
  Out-of-Order          : 0 (High:0)
  Duplicate ACKs        : 0
  RTT / Variance (High) : 0 ms (0 ms) / 0 ms (0 ms)

```

To display IPSec parameters on an IPSec-enabled tunnel:

```

switch:admin> portshow fciptunnel 24 -i

Tunnel: VE-Port:24 (idx:0, DP0)
=====
Oper State           : Enabled
TID                  : 24
Flags                : 0x00000000
IP-Extension         : Disabled
Compression          : None
QoS BW Ratio        : 50% / 30% / 20%
Fastwrite            : Disabled
Tape Pipelining     : Disabled
IPSec                : Enabled
IPSec-Policy         : myPolicy1
Load-Level (Cfg/Peer): Failover (Failover / Failover)
Local WWN            : 10:00:00:05:33:65:82:c8
Peer WWN             : 10:00:00:05:33:65:84:08
RemWWN (config)     : 00:00:00:00:00:00:00:00
cfgmask              : 0x0000001f 0x4000020c
Flow Status          : 0
ConCount/Duration   : 2 / 29m9s
Uptime               : 26s
Stats Duration      : 26s
Receiver Stats      : 28260 bytes / 164 pkts / 941.00 Bps Avg
Sender Stats        : 28096 bytes / 165 pkts / 935.00 Bps Avg
TCP Bytes In/Out    : 190584 / 236132
ReTx/000/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)   : 1 / 1 / 1 ms
Wan Util             : 0.0%
TxQ Util             : 0.0%

```

To display the FCIP tunnel hierarchy on the Brocade FX8-24:

```

switch:admin> portshow fciptunnel 1/12--hier
FCIP Tunnel 1/12
| High Level Tunnel Stats:
|   Operational Status: Up
|   Connected Count: 1
|   Max Comm Rate: 2500000
|   Compression Ratio: 1 : 1
|   Compressed Bytes: 0
|   Uncompressed Bytes: 0
|   Bytes In: 68942140
|     Bytes In Avg: 85
|   Bytes Out: 30983544
|     Bytes Out Avg: 16
|   Packets In: 92126
|     Packets In Avg: 0
|   Packets Out: 70728
|     Packets Out Avg: 0
| Aggregate TCP Stats:

```

```

TCP Bytes In: 842506236
TCP Bytes Out: 1446465008
TCP Packets In: 16072571
TCP Packets Out: 16063156
Retransmits: 0
Longest RTT: 0 ms
Out Of Sequence: 0
Slow Starts: 0
Circuit Count: 2
TCP Count: 24
Circuits:
--> Circuit 1/12.1
  | High Level Circuit Stats:
  | Operational Status: Up
  | Connected Count: 1
  | Max Comm Rate: 2500000
  | Bytes In: 712
  |   Bytes In Avg: 0
  | Bytes Out: 700
  |   Bytes Out Avg: 0
  | Packets In: 4
  |   Packets In Avg: 0
  | Packets Out: 5
  |   Packets Out Avg: 0
  | Aggregate TCP Stats:
  | TCP Bytes In: 68632640
  | TCP Bytes Out: 174210892
  | TCP Packets In: 2639463
  | TCP Packets Out: 2639467
  | Retransmits: 0
  | Longest RTT: 0 ms
  | Out Of Sequence: 0
  | Slow Starts: 0
  | TCP Count: 12
  | TCP Connections:
  | ---> TCP Conn 1/12.1:15240914
  | ---> TCP Conn 1/12.1:15240915
  | ---> TCP Conn 1/12.1:15240916
  | ---> TCP Conn 1/12.1:15241504
  | ---> TCP Conn 1/12.1:15241505
  | ---> TCP Conn 1/12.1:15241506
  | ---> TCP Conn 1/12.1:15241307
  | ---> TCP Conn 1/12.1:15241308
  | ---> TCP Conn 1/12.1:15241309
  | ---> TCP Conn 1/12.1:15241111
  | ---> TCP Conn 1/12.1:15241111
  | ---> TCP Conn 1/12.1:15241112
--> Circuit 1/12.0
  | High Level Circuit Stats:
  | Operational Status: Up
  | Connected Count: 1
  | Max Comm Rate: 2500000
  | Bytes In: 68941428
  |   Bytes In Avg: 82
  | Bytes Out: 30982844
  |   Bytes Out Avg: 11
  | Packets In: 92122
  |   Packets In Avg: 0
  | Packets Out: 70723
  |   Packets Out Avg: 0

```

```

| Aggregate TCP Stats:
|   TCP Bytes In: 773873596
|   TCP Bytes Out: 1272254116
|   TCP Packets In: 13433108
|   TCP Packets Out: 13423689
|   Retransmits: 0
|   Longest RTT: 0 ms
|   Out Of Sequence: 0
|   Slow Starts: 0
|   TCP Count: 12
|   TCP Connections:
|---> TCP Conn 1/12.0:15240899
|---> TCP Conn 1/12.0:15240900
|---> TCP Conn 1/12.0:15240901
|---> TCP Conn 1/12.0:15241489
|---> TCP Conn 1/12.0:15241490
|---> TCP Conn 1/12.0:15241491
|---> TCP Conn 1/12.0:15241292
|---> TCP Conn 1/12.0:15241293
|---> TCP Conn 1/12.0:15241294
|---> TCP Conn 1/12.0:15241096
|---> TCP Conn 1/12.0:15241097
|---> TCP Conn 1/12.0:15241097

```

To display a tunnel on a Brocade FX8-24 with automatic compression enabled:

```

switch:admin> portshow fciptunnel 1/21
Tunnel ID: 1/21
  Tunnel Description:
  Admin Status: Enabled
  Oper Status: Up
  Compression: On (Auto-Mode)
  Fastwrite: Off
  Tape Acceleration: Off
  TPerf Option: Off
  IPSec: Disabled
  QoS Percentages: High 50%, Med 30%, Low 20%
  Remote WWN: Not Configured
  Local WWN: 10:00:00:05:1e:52:fe:00
  Peer WWN: 10:00:00:05:1e:39:a4:76
  Circuit Count: 2
  Flags: 0x00000000
  FICON: Off

```

To display a tunnel on a Brocade FX8-24 with automatic compression enabled in summary view (The "A" flag indicates Auto-Mode):

```

switch:admin> portshow fciptunnel 1/21 -s
-----
Tunnel Circuit OpStatus Flags Uptime TxMBps RxMBps ConnCnt CommRt Met/G
-----
1/21 - Up A----- 31m29s 0.00 0.00 2 - -/-
1/21 0 1/xge1 Up ---4--s 31m29s 0.00 0.00 2 2500/2500 0/-
1/21 1 1/xge0 Up ---4-xs 31m29s 0.00 0.00 2 2500/2500 1/-
-----
Flags: tunnel: c=compression m=moderate compression

```

portShow

```

a=aggressive compression
A=Auto compression f=fastwrite t=Tapepipelining F=FICON
T=TPerf i=IPSec l=IPSec Legacy
Flags: circuit: s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6
L=Listener I=Initiator

```

To display a tunnel on a Brocade FX8-24 with IPsec enabled in legacy mode:

```

switch:admin> portshow fcipunnel 1/12
-----
Tunnel ID: 1/12
  Tunnel Description:
  Admin Status: Enabled
  Oper Status: Up
  Compression: Off
  Fastwrite: Off
  Tape Acceleration: Off
  TPerf Option: Off
  IPsec: Enabled (legacy)
  QoS Percentages: High 50%, Med 30%, Low 20%
  Remote WWN: Not Configured
  Local WWN: 10:00:00:05:1e:52:fe:00
  Peer WWN: 10:00:00:05:1e:39:a4:76
  Circuit Count: 1
  Flags: 0x00000000
  FICON: Off

```

To display a summary view of the FCIP tunnel and circuits showing then legacy flag for the tunnel:

```

switch:admin> portshow fcipunnel 8/12 -s
-----
Tunnel Circuit OpStatus Flags Uptime TxMBps RxMBps ConnCnt CommRt Met/G
-----
8/12 -          InProg  cft--l- 3m3s  0.00   0.00    1    -    -
8/12 0 8/xge1 InProg  ---4v-s 3m3s  0.00   0.00    1  5000/5000 0/-
8/12 1 8/xge0 InProg  ---4-xs 3m0s  0.00   0.00    1  1000/1000 0/-
8/12 2 8/xge1 InProg  ---6--s 3m0s  0.00   0.00    1  2000/2000 0/-
-----
Flags: tunnel: c=compression m=moderate compression
          a=aggressive compression
          A=Auto compression f=fastwrite t=Tapepipelining F=FICON
          T=TPerf i=IPSec l=IPSec Legacy
Flags: circuit: s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6

```

To display a summary view of an FCIP tunnel and its circuits on the Brocade FX8-24 with the **--perf** option :

```

switch:admin> portshow fcipunnel 7/12 --summary --perf
-----
Tunnel Circuit TxMBps RxMBps ComRatio RTT(ms) ReTx OutOfSeq SlowStart
-----
1/22 -          0.00   0.00  1.00:1    -      0      0      0
1/22 0 1/xge0  0.00   0.00    -        0      0      0
1/22 1 1/xge1  0.00   0.00    -        0      0      0
-----

```

To display a summary view of the QoS configuration on FCIP tunnels and circuits on the Brocade FX8-24:

```
switch:admin> portshow fciptunnel 7/12 --summary --qos
-----
Tunnel  Circuit  Priority  Uptime    Tx Bps    Rx Bps  TxPkts/s  RxPkts/s
-----
1/22    -           -        2d23h10m  0.00      0.00    0.00      0.00
1/22    -           F-Class  2d23h10m  0.00      0.00    0.00      0.00
1/22    -           High     2d23h10m  0.00      0.00    0.00      0.00
1/22    -           Medium   2d23h10m  0.00      0.00    0.00      0.00
1/22    -           Low      2d23h10m  0.00      0.00    0.00      0.00
-----
```

```
switch:admin> portshow fciptunnel 7/12 --summary --qos --circuit
-----
Tunnel  Circuit  Priority  Uptime    Tx Bps    Rx Bps  TxPkts/s  RxPkts/s
-----
1/22    -           -        2d23h11m  0.00      0.00    0.00      0.00
1/22    -           F-Class  2d23h11m  0.00      0.00    0.00      0.00
1/22    -           High     2d23h11m  0.00      0.00    0.00      0.00
1/22    -           Medium   2d23h11m  0.00      0.00    0.00      0.00
1/22    -           Low      2d23h11m  0.00      0.00    0.00      0.00
1/22    0 1/xge0  -        2d23h11m  0.00      0.00    0.00      0.00
1/22    0 1/xge0  F-Class  2d23h11m  0.00      0.00    0.00      0.00
1/22    0 1/xge0  Medium   2d23h11m  0.00      0.00    0.00      0.00
1/22    0 1/xge0  Low      2d23h11m  0.00      0.00    0.00      0.00
1/22    1 1/xge1  -        2d23h11m  0.00      0.00    0.00      0.00
1/22    1 1/xge1  F-Class  2d23h11m  0.00      0.00    0.00      0.00
1/22    1 1/xge1  High     2d23h11m  0.00      0.00    0.00      0.00
1/22    1 1/xge1  Medium   2d23h11m  0.00      0.00    0.00      0.00
1/22    1 1/xge1  Low      2d23h11m  0.00      0.00    0.00      0.00
-----
```

To display the IP addresses configured for the circuits on the Brocade FX8-24:

```
switch:admin> portshow fciptunnel --circuit --config
-----
Tunnel  Circuit  AdminSt  Flags    Local IP                Remote IP
-----
8/12    -           Enabled  -----
8/12    0 8/xge1  Enabled  ---4--s  192.168.0.16            192.168.0.116
8/24    -           Enabled  -----
8/24    0 8/xge0  Enabled  ---4--s  192.168.0.11            192.168.0.111
-----
Flags (tunnel): c=compression m=moderate compression a=aggressive compression
                A=Auto compression f=fastwrite t=Tapipelining F=FICON
                T=TPerf i=IPSec l=IPSec Legacy
(circuit):     s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6
                L=Listener I=Initiator
```

To display all FCIP circuits on the Brocade FX8-24:

```
switch:admin> portshow fcipcircuit all
-----
Tunnel  Circuit  OpStatus  Flags    Uptime  TxMBps  RxMBps  ConnCnt  CommRt  Met/G
-----
1/12    0 1/xge1  Up       ---4--s  3d2m    000     0.00    1  2500/2500  0/-
1/12    1 1/xge0  Up       ---4-xs  3d2m    0.00    0.00    1  2500/2500  1/-
1/21    0 1/xge1  Up       ---4--s  3d1m    0.00    0.00    1  2500/2500  0/-
```

portShow

```
1/21 1 1/xge0 Up ---4-xs 3d2m 0.00 0.00 1 2500/2500 1/-
1/22 0 1/xge0 Up ---4--s 3d1m 0.00 0.00 1 2500/2500 0/-
1/22 1 1/xge1 Up ---4-xs 3d2m 0.00 0.00 1 2500/2500 1/-
1/31 0 1/xge0 Up ---4--s 3d2m 0.00 0.00 1 2500/2500 0/-
1/31 1 1/xge1 Up ---4-xs 3d2m 0.00 0.00 1 2500/2500 1/-
```

```
-----
Flags:circuit:s=sack v=VLAN Tagged x=crossport 4=IPv4 6=IPv6
      T=Test(CPerf) L=Listener I=Initiator
```

To display the details for all FCIP circuits on the Brocade FX8-24 (Note that both circuits are configured as initiators):

```
switch:admin> portshow fcipcircuit all --detail
```

```
-----
Circuit ID: 1/12.0
  Circuit Num: 0
  Admin Status: Enabled
  Oper Status: Up
  Connection Type: Initiator
  Remote IP: 192.168.12.200
  Local IP: 192.168.12.100
  Metric: 0
  Failover Group ID: (Not Config/Active)
  Min Comm Rt: 2500000
  Max Comm Rt: 2500000
  SACK: On
  Min Retrans Time: 100
  Max Retransmits: 8
  Keepalive Timeout: 10000
  Path MTU Disc: 0
  VLAN ID: (Not Configured)
  L2CoS: (VLAN Not Configured)
  DSCP: F: 0 H: 0 M: 0 L: 0
  Flags: 0x00000000
```

```
-----
Circuit ID: 1/12.1
  Circuit Num: 1
  Admin Status: Enabled
  Oper Status: Up
  Connection Type: Initiator
  Remote IP: 192.168.12.201
  Local IP: 192.168.12.101
  Metric: 1
  Failover Group ID: (Not Config/Active)
  Min Comm Rt: 2500000
  Max Comm Rt: 2500000
  SACK: On
  Min Retrans Time: 100
  Max Retransmits: 8
  Keepalive Timeout: 10000
  Path MTU Disc: 0
  VLAN ID: (Not Configured)
  L2CoS: (VLAN Not Configured)
  DSCP: F: 0 H: 0 M: 0 L: 0
```



```

      Flags: 0x00000000
-----
Circuit ID: 1/21.0
  Circuit Num: 0
  Admin Status: Enabled
(Output truncated)

```

To display additional performance parameters for a circuit:

```

switch:admin> portshow fcipcircuit 1/12 0 --perf
-----
Tunnel ID: 1/12
  Tunnel Description:
  Admin Status: Enabled
  Oper Status: Up
  Compression: Off
  Fastwrite: Off
  Tape Acceleration: Off
  TPerf Option: Off
  IPSec: Disabled
  QoS Percentages: High 50%, Med 30%, Low 20%
  Remote WWN: Not Configured
  Local WWN: 10:00:00:05:1e:52:fe:00
  Peer WWN: 10:00:00:05:1e:39:a4:76
  Circuit Count: 2
  Flags: 0x00000000
  FICON: Off
  Oper Status: Up
  Flow Ctrl State: Off
  Connected Count: 1
  Tunnel Duration: 3 days, 19 hours, 54 minutes, 5 seconds
  Compression Statistics:
    0 Uncompressed Bytes
    0 Compressed Bytes
    1.00 : 1 Compression Ratio
  Performance Statistics: Overall Throughput
    31073824 Output Bytes
    16 Bps 30s Avg, 93 Bps Lifetime Avg
    70932 Output Packets
    0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
    69170308 Input Bytes
    85 Bps 30s Avg, 209 Bps Lifetime Avg
    92393 Input Packets
    0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
  TCP Stats:
    1450493536 Output Bytes
    16107875 Output Packets
    844887492 Input Bytes
    16117320 Input Packets
    Retransmits: 0
    Round Trip Time: 0 ms
    Out Of Order: 0
    Slow Starts: 0

```

To display QoS prioritization for the default circuit:

```
switch:admin> portshow fcipcircuit 1/12 0 --perf --qos
-----
Circuit ID: 1/12.0
  Circuit Num: 0
  Admin Status: Enabled
  Oper Status: Up
  Connection Type: Default
  Remote IP: 192.168.12.200
  Local IP: 192.168.12.100
  Metric: 0
  Failover Group ID: (Not Config/Active)
  Min Comm Rt: 2500000
  Max Comm Rt: 2500000
  SACK: On
  Min Retrans Time: 100
  Max Retransmits: 8
  Keepalive Timeout: 10000
  Path MTU Disc: 0
  VLAN ID: (Not Configured)
  L2CoS: (VLAN Not Configured)
  DSCP: F: 0 H: 0 M: 0 L: 0
  Flags: 0x00000000
  Flow Ctrl State: Off
  Connected Count: 1
  Circuit Duration: 3 days, 19 hours, 57 minutes, 42 seconds
  Performance Statistics - Priority: F-Class
    Oper Status: Up
    Flow Ctrl State: Off
    Connected Count: 1
    Duration: 3 days, 19 hours, 57 minutes, 42 seconds
    4732308 Output Bytes
      10 Bps 30s Avg, 14 Bps Lifetime Avg
    41359 Output Packets
      0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
    7223112 Input Bytes
      27 Bps 30s Avg, 21 Bps Lifetime Avg
    45277 Input Packets
      0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
    TCP Stats:
      318824216 Output Bytes
      3419555 Output Packets
      184935644 Input Bytes
      3419605 Input Packets
      Retransmits: 0
      Round Trip Time: 0 ms
      Out Of Order: 0
      Slow Starts: 0
  Performance Statistics - Priority: High
    Oper Status: Up
    Flow Ctrl State: Off
    Connected Count: 1
    Duration: 3 days, 19 hours, 57 minutes, 41 seconds
    0 Output Bytes
      0 Bps 30s Avg, 0 Bps Lifetime Avg
    0 Output Packets
      0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
    0 Input Bytes
      0 Bps 30s Avg, 0 Bps Lifetime Avg
```

```

0 Input Packets
  0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
TCP Stats:
  308550696 Output Bytes
  3310618 Output Packets
  176126080 Input Bytes
  3310620 Input Packets
  Retransmits: 0
  Round Trip Time: 0 ms
  Out Of Order: 0
  Slow Starts: 0
Performance Statistics - Priority: Medium
  Oper Status: Up
  Flow Ctrl State: Off
  Connected Count: 1
  Duration: 3 days, 19 hours, 57 minutes, 43 seconds
  26358236 Output Bytes
    17 Bps 30s Avg, 79 Bps Lifetime Avg
  29611 Output Packets
    0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
  61983140 Input Bytes
    298 Bps 30s Avg, 187 Bps Lifetime Avg
  47166 Input Packets
    0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
TCP Stats:
  340704236 Output Bytes
  3429074 Output Packets
  239375444 Input Bytes
  3438474 Input Packets
  Retransmits: 0
  Round Trip Time: 0 ms
  Out Of Order: 0
  Slow Starts: 0
Performance Statistics - Priority: Low
  Oper Status: Up
  Flow Ctrl State: Off
  Connected Count: 1
  Duration: 3 days, 19 hours, 57 minutes, 42 seconds
  0 Output Bytes
    0 Bps 30s Avg, 0 Bps Lifetime Avg
  0 Output Packets
    0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
  0 Input Bytes
    0 Bps 30s Avg, 0 Bps Lifetime Avg
  0 Input Packets
    0 pkt/s 30s Avg, 0 pkt/s Lifetime Avg
TCP Stats:
  308551588 Output Bytes
  3310627 Output Packets
  176126612 Input Bytes
  3310630 Input Packets
  Retransmits: 0
  Round Trip Time: 0 ms
  Out Of Order: 0
  Slow Starts: 0

```

To display a summary view for the circuits on the Brocade 7840:

```
switch:admin> portshow fciptunnel --circuit
```

| Tunnel | Circuit | OpStatus | Flags | Uptime | TxMBps | RxMBps | ConnCnt | CommRt | Met/G |
|--------|---------|----------|----------|--------|--------|--------|---------|-----------|-------|
| 24 | - | Up | ----- | 36s | 0.00 | 0.00 | 1 | - | -/- |
| 24 | 0 ge2 | Up | ---ah--4 | 36s | 0.00 | 0.00 | 1 | 2500/5000 | 0/- |
| 24 | 1 ge3 | Up | ---ah--4 | 10s | 0.00 | 0.00 | 1 | 2500/5000 | 0/- |
| 28 | - | Up | ----- | 13s | 0.00 | 0.00 | 1 | - | -/- |
| 28 | 0 ge4 | Up | ---a---4 | 13s | 0.00 | 0.00 | 1 | 2500/5000 | 0/- |
| 34 | - | InProg | ----- | 0s | 0.00 | 0.00 | 0 | - | -/- |
| 34 | 0 ge2 | InProg | ---ah--4 | 0s | 0.00 | 0.00 | 0 | 2500/5000 | 0/- |
| 34 | 1 ge3 | InProg | ---ah--4 | 0s | 0.00 | 0.00 | 0 | 2500/5000 | 0/- |

Flags (tunnel): i=IPSec f=Fastwrite T=TapePipelining F=FICON r=ReservedBW
a=FastDeflate d=Deflate D=AggrDeflate
(circuit): h=HA-Configured v=VLAN-Tagged p=PMTU 4=IPv4 6=IPv6
ARL a=Auto r=Reset s=StepDown t=TimedStepDown

To display the IP addresses configured for the circuits on the Brocade 7840:

```
switch:admin> portshow fciptunnel 24 --circuit
```

```
Tunnel: VE-Port:24 (idx:0)
```

```
=====
Oper State           : In Progress
TID                  : 24
Flags                 : 0x00000000
Compression          : None
QoS BW Ratio         : 50% / 30% / 20%
Fastwrite            : Disabled
Tape Pipelining      : Disabled
IPSec                 : Disabled
Local WWN             : 10:00:00:05:1e:65:7d:08
Peer WWN              : 00:00:00:00:00:00:00:00
RemWWN (config)      : 00:00:00:00:00:00:00:00
cfgmask              : 0x40000008 0x0000001f
Failover Count       : 0
Flow Status          : 0
ConCount/Duration    : 0 / 0s
Receiver Stats       : 0 bytes / 0 pkts / 0.00 Bps Avg
Sender Stats         : 0 bytes / 0 pkts / 0.00 Bps Avg
TCP Bytes In/Out     : 0 / 0
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)    : 0 / 0 / 0 ms
```

```
Circuit 24.0 (dp:0)
```

```
=====
Admin/Oper State     : Enabled / In Progress
Flags                 : 0x00000000
IP Addr (L/R)        : 192.168.2.20 ge2 - 192.168.2.10
HA IP Addr (L/R)     : 192.168.2.21 ge2 - 192.168.2.11
Configured Comm Rates: 2500000 / 5000000 kbps
Peer Comm Rates      : 0 / 0 kbps
Actual Comm Rates    : 2500000 / 5000000 kbps
Keepalive Timeout    : 6000 ms
Metric                : 0
Connection Type      : Default
ARL-Type              : Auto
PMTU                  : Disabled
```

```

Failover Group      : 0
VLAN-ID            : NONE
L2Cos (f/h/m/l)    : 0/0/0/0
DSCP (f/h/m/l)     : 0/0/0/0
cfgmask            : 0x00003c2f 0x40000000
Flow Status        : 1
ConCount/Duration  : 0 / 0s
Receiver Stats     : 0 bytes / 0 pkts / 0.00 Bps Avg
Sender Stats       : 0 bytes / 0 pkts / 0.00 Bps Avg
TCP Bytes In/Out   : 0 / 0
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)  : 0 / 0 / 0 ms

```

Circuit 24.1 (dp:0)

```

=====
Admin/Oper State   : Enabled / In Progress
Flags              : 0x00000000
IP Addr (L/R)      : 192.168.3.20 ge3 - 192.168.3.10
HA IP Addr (L/R)   : 192.168.3.21 ge3 - 192.168.3.11
Configured Comm Rates: 2500000 / 5000000 kbps
Peer Comm Rates    : 0 / 0 kbps
Actual Comm Rates  : 2500000 / 5000000 kbps
Keepalive Timeout  : 6000 ms
Metric             : 0
Connection Type    : Default
ARL-Type           : Auto
PMTU               : Disabled
Failover Group     : 0
VLAN-ID            : NONE
L2Cos (f/h/m/l)    : 0/0/0/0
DSCP (f/h/m/l)     : 0/0/0/0
cfgmask            : 0x00003c2f 0x40000000
Flow Status        : 1
ConCount/Duration  : 0 / 0s
Receiver Stats     : 0 bytes / 0 pkts / 0.00 Bps Avg
Sender Stats       : 0 bytes / 0 pkts / 0.00 Bps Avg
TCP Bytes In/Out   : 0 / 0
ReTx/OOO/SloSt/DupAck: 0 / 0 / 0 / 0
RTT (min/avg/max)  : 0 / 0 / 0 ms

```

To display FCP emulation statistics using the command with two levels and a single command:

```
switch:admin> portshow xtun 1/13 -fcp -itl -itn -stats
```

To display FCP emulation statistics, information, and configuration information using the command with a single level and three commands:

```
switch:admin> portshow xtun 1/13 -fcp -itl -stats \
-info -cfg
```

To display FICON statistics on the Brocade SX6 blade:

```
switch:admin> portshow xtun 8/27 -ficon -stats
```

```

FICON FCIP VePort=219 Statistics:
Emulation Common Statistical Counts (decimal)
=====
TotalIngressFrames = 307196261

```

portShow

```
TotalEgressFrames      =          21636019
TotalFCEgressFrames    =          10323404
TotalCmds               =           7420486
TotalEmulDvcLvlAcks    =              734
TotalEmulatedOps       =          1889025
Idle Accepted Status   =              92
Idle Not Acpt Status   =              14
Idle Accepted ATN      =              4
MaxRetryQueueDepth    =              0   MaxEgressQueueDepth =          3360
TotalCUBusyResponses  =              5   TotalCUEndResponses =              5
TotalEmulatedCUBusys  =              5   TotalEmulCUEnd      =              5
WIRE Buffer Percentages LocalFree= 99 LocalLow= 99 PeerFree= 99 PeerLow=99
XBAR Buffer Percentages LocalFree=100 LocalLow=100 PeerFree=100 PeerLow=100
```

FICON Configuration Status:

```
-----
FICON Host side Paths =          0 Device Side Paths=          0
FICON LPARS Connected =          5   Curr LCUs =          12 Total Devices= 128
Current Egress Q Count=        3360   RetryQCnt =          0
Current Free Headers  =         348   HdrsInUse =          0
Active Emulation Count= 12
```

No Emulated XRC Operations

Tape Write Emulation Statistical Counts (decimal)

```
-----
Current Bytes in write pipe      =          0
Maximum Bytes in write pipe      =          14680064
Largest write chain processed    =          229377
Total number of emulated Write Bytes=        249337825280
Number of emulated Write Chains  =          1086924
Total number of emulated Write CCWs =          4348063
Average Emulated Writes Blocksize =          57344
Average Writes in Emulated Chains =              4
Write emulation slowdowns        =          30980
Slowdowns at Start of Chain      =          407
Slowdowns at End of chain        =          5291
Current Host side Write FDCB Count =          0
Max Concurrent Write FDCB Count  =          8
Current Write Limited FDCB Count =          0
```

Tape Read Emulation Statistical Counts (decimal)

```
-----
Total number of emulated Read Bytes =          73229283205
Total number of emulated Read Chains=          802101
Total number of emulated Read CCWs  =          1180476
Average Emulated Bytes per chain    =          91296
Average Emulated Read Blocksize     =          62033
Average CCWs in Emulated Chains     =              1
Current Read FDCBs Count            =          8
Max Concurrent Read FDCB Count      =          0
Current Read Limited FDCB Count     =          8
```

No Emulated Tera Operations

No Emulated Printer Operations

FICON Debug Flags (ftrace and others) = 0x77c90000 (Default = 0x77c90000)

```

-----
Bit 31 [0x80000000] = 0 : TRIGGER_ON_SELRESET
Bit 30 [0x40000000] = 1 : TRIGGER_ON_PURGEPATH
Bit 29 [0x20000000] = 1 : TRIGGER_ON_RRS_MISS
Bit 28 [0x10000000] = 1 : TRIGGER_ON_LRJ

Bit 27 [0x08000000] = 0 : TRIGGER_ON_UNIT_CHECK
Bit 26 [0x04000000] = 1 : TRIGGER_ON_LOOKUP_FAIL
Bit 25 [0x02000000] = 1 : TRIGGER_ON_FDCB_ABORT
Bit 24 [0x01000000] = 1 : TRIGGER_ON_NOFDCB_ABORT

Bit 23 [0x00800000] = 1 : TRIGGER_ON_LINKDOWN
Bit 22 [0x00400000] = 1 : TRIGGER_ON_SENSE_CMD
Bit 21 [0x00200000] = 0 : TRIGGER_ON_BUSYATTN
Bit 20 [0x00100000] = 0 : TRIGGER_ON_XRCUNS

Bit 19 [0x00080000] = 1 : DISPLAY_FDCB_ON_ERROR
Bit 18 [0x00040000] = 0 : TRIGGER_ON_LOW_WIREPOOL
Bit 17 [0x00020000] = 0 : TRIGGER_ON_READ_UNITCHECK
Bit 16 [0x00010000] = 1 : TRIGGER_ON_STATE_SAVE

Bit 15 [0x00008000] = 0 : TRIGGER_ON_SUSPEND
Bit 14 [0x00004000] = 0 : MINIMIZE_RASLOGS
Bit 13 [0x00002000] = 0 : TRIGGER_ON_LONG_IO
Bit 12 [0x00001000] = 0 : FORCE_DISABLE_IDLE_STACC

Bit 11 [0x00000800] = 0 : FORCE_DISABLE_3203_PRINT
Bit 10 [0x00000400] = 0 : FORCE_DISABLE_3211_PRINT
Bit 9 [0x00000200] = 0 : FORCE_DISABLE_3800_PRINT
Bit 8 [0x00000100] = 0 : FORCE_DISABLE_3900_PRINT

Bit 5 [0x00000020] = 0 : DISABLE_TAPE_MULTIPATH_PROTECTION
Bit 4 [0x00000010] = 0 : ACCEPT_READTRACK_STATUS

Bit 3 [0x00000008] = 0 : FORCE_DISABLE_READ_PIPE
Bit 2 [0x00000004] = 0 : FORCE_DISABLE_WRITE_PIPE
Bit 1 [0x00000002] = 0 : FORCE_DISABLE_XRC_EMUL
Bit 0 [0x00000001] = 0 : DISABLE_PERSIST_IU_PACE

```

To display the FICON Device Path Blocks:

```
switch:admin> portshow xtun 23 -ficon -fdpb
```

```

FDPB (FICON Device Path Block - one per path) Count = 4
----- \
(0x)      Side Path:          Emul Type Tag      \
=====  ====  =====  =====  =====  \
041004E000  H  0x1763016401***** Yes  Tape 0x14340000 \
  flg=10 parms=2001F0 typ=002086 mod=A04 mfg=IBM plnt=02 \

0410062880  H  0x17630164FE***** No   FCUP 0x14FD0000 \
0410052000  H  0x1763086403***** Yes  Disk 0x14830000 \
  flg=10 parms=2001F2 typ=002086 mod=A04 mfg=IBM plnt=02 \

0410084000  H  0x176308640E***** Yes  ESCN 0x14A40000 \
  flg=10 parms=2001F2 typ=002086 mod=A04 mfg=IBM plnt=02 \

```

portShow

```

-----
  EgrOx IngOx Valid VTN
  =====
    0     1   Y     1
  sq=0000000E77DF tg=80F0

    0     0   Y     1
    0     0   Y     1
  sq=0000000E77DF tg=80F2

    1     1   Y     1
  sq=0000000E77DF tg=80F2
  (Output split)

```

To display the FICON Channel Blocks (all blocks and a specified block):

```

switch:admin> portshow xtun 16 -ficon -fchb 041055B680

FCHB (FICON CHannel Block - one per LPAR) Count = 1
-----
(0x)          Side Path:                CU Count  Emul ChTIN CuTIN
=====
041055B680    D 0x106301640106****    0x0004    0x0000 0x0000

FCHB Flags:
  tinInProgress=N   emuTinAckPending=N   emuTirInProgress=N
  emuTirAckPending=N emuTinLackPending=N   emuTirPending=N
  emuTirReceived=N emuTinFinalAckPend=N   emuTirPending=N
  emuTinSuccessful=N emuTinReceived=N     emuTirSent=N
  fcrPresent=N      fchbValid=Y

FCHB Control variables:
  tinOrigOxid=0xFFFF  tirOrigOxid=0xFFFF  tinAckRxid=0xFFFF
  tinAckOxid=0xFFFF  allocChOxid=0xFFFF  allocChOIdx=0x0000
  tinTirOxid=0xFFFF

```

To display FICON FCUB information (FICON Images accessed through the tunnel)

```

switch:admin> portshow xtun 23 -ficon -fcub

FCUB (FICON Control Unit Block) Count = 29

FC Egress parms:
(0x)          Side Path:                devRange DevCnt Model \
=====
0410083800    H 0x17630164010604**    00-0F    0x10  3490 \
0410047500    H 0x17630164010605**    00-0F    0x10  3490 \
0410048E80    H 0x17630164010606**    00-0F    0x10  3490 \
0410049700    H 0x17630164010607**    00-0F    0x10  3490 \
0410054000    H 0x17630164010004**      -         0x00  0000 \
041006B100    H 0x17630864030600**    00-7F    0x80  3990 \
041006C000    H 0x17630864030601**    00-7F    0x80  3990 \
04100A4000    H 0x17630864030000**      -         0x00  0000 \
0410072000    H 0x176308640E0600**      -         0x00  0000 \
0410073100    H 0x176308640E0601**      -         0x00  0000 \
0410075100    H 0x176308640E0602**      -         0x00  0000 \

```



```

0410076880   H   0x176308640E0603**   -   0x00   0000   \
0410078000   H   0x176308640E0604**   -   0x00   0000   \
0410078880   H   0x176308640E0605**   -   0x00   0000   \
041007A000   H   0x176308640E0606**   -   0x00   0000   \

```

```

FC Egress parms:
Type  LPE BL CHPID          Tag V Ve P B P VC EGID
====  == == ===== = == = = = == =====
TAPE  Yes  00F0 0x15698E1E Y  Y 2 0 1 02 000D
TAPE  Yes  00F0 0x15B5E15F Y  Y 2 0 2 02 000D
TAPE  Yes  00F0 0x15D1509C Y  Y 2 0 4 02 000D
TAPE  Yes  00F0 0x150D3FDD Y  Y 2 0 2 02 000D
UNKN  No   00F0 0x15CDB71F Y  Y 2 0 1 02 000D
DISK  Yes  00F2 0x1538206F Y  Y 2 0 4 02 0007
DISK  Yes  00F2 0x15E44F2E Y  Y 2 0 1 02 0007
UNKN  No   00F2 0x159C196E Y  Y 2 0 3 02 0007
UNKN  No   00F2 0x152D3BA3 Y  Y 2 0 3 02 0007
UNKN  No   00F2 0x15F154E2 Y  Y 2 0 4 02 0007
UNKN  No   00F2 0x1595E521 Y  Y 2 0 4 02 0007
UNKN  No   00F2 0x15498A60 Y  Y 2 0 3 02 0007
UNKN  No   00F2 0x155C86A7 Y  Y 2 0 1 02 0007
UNKN  No   00F2 0x1580E9E6 Y  Y 2 0 1 02 0007
UNKN  No   00F2 0x15E45825 Y  Y 2 0 3 02 0007
(output split and truncated)

```

To reset the Device Path Block statistics:

```
switch:admin> portshow xtun 16 -ficon -fdpb -clear
```

To display a FICON Teradata, Tape, or XRC emulation statistics for a specified Device Control Block:

```

switch:admin> portshow xtun 16 -ficon -fdcb 0x041008B980
FDCB (FICON Device Control Block):Port=10 Side=Host Active=No
DeviceType=TERA
-----
FDCB
(0x)      hDom hPrt dDom dPrt lch  lcu  dev  state
=====  ==  ==  ==  ==  ==  ==  ==  ==
0x041008B980 63  04  64  02  02  08  03  0x00

CONTROL
=====
active           = No Send_SYR           = No crrSet           = No
xrcEstablished  = No sssSet             = No dvcAckEmulInProg = No
discardIgrFrames= No resetEmulPending    = No interceptLack   = No
interceptLack2  = No onStartPendingQueue= No discardEgrFrames = No
deferDeviceType = No emulBaAcceptPending= No mappingOxid      = No
xrcEmulEnable   = Yes tapeWriteEmulEnable= Yes tapeReadEmulEnabled= Yes
dvcAckEmulEnable= Yes fdcbLocked         = Yes vtnValid        = Yes
fcrPresent      = No deviceNotInstalled = No onCuBusyQueue    = No
egressParmsSet  = Yes statusFlags       = 0x4000

QUEUE COUNTS
=====
fcEgressQueue   =          0  egressQMax =          0
ficonRetryQueue =          0  retryQMax  =          0
cmdHdrQueue     =          0  cmdHdrQMax =         32

```

```

MISC (hex)
=====
state      = 0x00      prevState    = 0x18      lastStateArray= 0x1C371415
statusFlags = 0x4000
errorCode  = 0x00      ingressOxid  = 0xFFFF
egressOxid = 0xFFFF    allocatedOxid= 0xFFFF    unsolIngrOxid = 0xFFFF
lastStatus = 0x0C      lastCmd      = 0x00      tokenFlags    = 0x00
lastCmdArray= 0x773E3E64 lastStsArray= 0x08040C0C1 \
                                           astXprtArray=0x08040C0C
lastSeqId  = 0x00      curSeqId     = 0xB1
lastXportEmulMsg = 0x4F  statusFctlHi = 0x18      curIuCnt      = 0x0001
token      = 0x00000000
endingChOxid = 0xFFFF  endingCuOxid= 0xFFFF
emulDvcAckSeqs= 0x00000003 lastEvtFromCu= 0x00  priorEvtFrmCu = 0x85
abortOxid  = 0xFFFF  fdcbWorkSched = 0
deviceType = 0x54455241

```

```

EMUL_HDR (hex)
=====
funct = 0x00 msg = 0x00
parm1 = 0x00 parm2 = 0x0000 parm3 = 0x00000000

```

```

STATS (dec)
=====
cmdCount      =      274865
emulatedOps   =      274565  resequencedIuCount =      0
cancelCount   =      0      selectiveResetCount =      1
purgePathCount =      0      abortCount          =      0
normStatus    =      274408  attnStatus          =      0
attnBusyStatus =      0      ceStatus            =      0
deStatus      =      0      retryStatus         =      0
immRetryStatus =      0      devBusyStatus       =      0
cuBusyStatus  =      0      ceDeUxStatus        =      1
deUxStatus    =      0      pendDeStatus        =      0
deUcStatus    =      0      unusualStatus       =      0
chLinkBusyCount =      1      cuLinkBusyCount     =      0
unusualStsArray = 0x00000000
pGb           = 0x0000000000
tag           = 0x0285CF9E
senseDeviceType = 0xFF30884000000000
ingressFrameCount= 0x000000000011ADCE
egressFrameCount = 0x00000000000432AB

```

```

TERADATA CONTROL
=====
writeSuspendFlag = OFF  finalStatusPndg = OFF  waitCuRsp      = OFF
cuBusyRsp        = OFF  suspendCompleted= OFF  rexmitPending  = OFF
synStatusSent    = OFF  unitCheckPndg  = OFF  unAnsweredExchg = OFF
queuedCntlFrame  = OFF  controlFunctAc = OFF  controlFunctDon = OFF
unitChkPresented = OFF  senseDataSaved = OFF  snsDataPresente = OFF
waitingForSenseData= OFF  lackOwed       = OFF  rexmitLbySent   = OFF
readEmulAct      = OFF  writeEmulAct    = OFF  iuPacingNeeded  = ON

```

```

TERADATA MISC
=====
writeTrigger      = 0x0000      readBlkTrigger  = 0x0000
teraDhEndCount   = 0x0
teraEsCount       = 0x0000      writeOpsInPipe  = 0x00000000
maxWritePipe     = 0x0000      maxReadPipe     = 0x0002
teraCmdCode      = 0x64         teraStatus      = 0x0d

```

```

emulTeraWriteOps      = 0x00043085      emulTeraReadOps = 0x00000000
emulTeraWrtBytes     = 0x00000001679cb8b0
emulTeraReadBytes    = 0x0000000000000000
emulTeraWriteCcws    = 0x00043085      emulTeraReadCcws= 0x00000000
multUnitCheckCnt     = 0x00000000
currentWriteChain    = 160
largestWriteChain    = 43936
bytesInWritePipe     = 0
atMaxWrtBytesCount  = 0
IU Pacing Values:crrsIssued=0 creditsOutstanding=0 currentCreditBurst=0

```

Historic Emulation Headers:

```

=====
current index = 1
ToPeerEmulHdr[0]=funct=2:msg=1:p1=0x00:p2=0x0000:p3=0x0285cf9e
ToPeerEmulHdr[1]=funct=2:msg=7:p1=0x00:p2=0x3085:p3=0x00000000
ToPeerEmulHdr[2]=funct=2:msg=7:p1=0x00:p2=0x3085:p3=0x00000000
ToPeerEmulHdr[3]=funct=2:msg=18:p1=0x00:p2=0x3085:p3=0x00000000
current index = 0
FromPeerEmulHdr[0]=funct=2:msg=15:p1=0x0c:p2=0x3066:p3=0x00000003
FromPeerEmulHdr[1]=funct=7:msg=8:p1=0x5f:p2=0x0062:p3=0x00000000
FromPeerEmulHdr[2]=funct=7:msg=6:p1=0x00:p2=0x0000:p3=0x00000000
FromPeerEmulHdr[3]=funct=7:msg=8:p1=0x5f:p2=0x0062:p3=0x00000000
convIngressOxid      = 0xffff      convIngressOxid  = 0xffff
timedOps              = 0x00043066

```

TERA IO TIMING:

```

=====
last elapsedTime     = 0.397
maxElapsedTime       = 0.504
totalTime            = 145.132
longIOtimes          = 0x0000005D
startOfChainDelayCnt= 0x0000      endOfChainDelayCnt = 0x5B
readPacedCount       = 0x0000      writePacedCount    = 0x0000

```

Ave Chain SIO time = 0.000 seconds

ReadPipe Cntls:

```

=====
readBlkEmulation=OFF      entireCP=OFF      waitForAccept=OFF
      endOfCP=OFF      commandRetry=OFF  goToIdleState=OFF
      sentUExceptn=OFF    earlyEnd=OFF      dackPending=OFF
      sofCNRreceived=OFF  noActiveIO=OFF    sendSOFRep=OFF
      blkIneligible=OFF   rdBlkValid1=OFF   rdBlkValid2=OFF
earlyIdleStatus=ON      reducePipe=OFF    doBSNoOp=OFF
      waitForDe=OFF      readCpLimited=OFF
statusFramesOnQueue = 0x0000      readBlocksSent   = 0x0000
readBlocksRequested = 0x0000      readCpsRequested = 0x0000

```

To display FICON emulation statistics on the Brocade SX6 (as the following examples show, this command can display emulation statistics for Teradata, Tape, XRC, and Printer depending on the FICON devices and FICON Emulation Features that are enabled on the tunnel.):

```
switch:admin> portshow xtun 8/27 -ficon -emultera
                                TAPE EMULATION STATS
+-----+-----+-----+-----+-----+-----+
|      FDCB Ptr      |      Path      | H|State|Emul|Emul|Rtry| Emulated |
| (0x)              | (0x)           | D|     |Pipe|Q'd | Qd | Tape Ops |
+-----+-----+-----+-----+-----+-----+
|0x8000000500EBC900|DB6841614D02030E|D| 0x3C|0x11|0007|0000|      149487|
|0x8000000500EC3400|DB6841614D020401|D| 0x3C|0x11|0005|0000|      149436|
+-----+-----+-----+-----+-----+-----+
Emulated |RdAvg |Emulated |WtAvg |
Read CCWs| Size |Write CCWs| Size |
+-----+-----+-----+-----+
      149487| 65534|          0|      0|
      149436| 65533|          0|      0|
(Output split and truncated)
```

```
switch:admin> portshow xtun 23 -ficon -emultape
                                TAPE EMULATION STATS
+-----+-----+-----+-----+-----+-----+
|      FDCB Ptr      |      Path      | H|State|Cmds| Cmd|Data|Data| Emulated |
| (0x)              | (0x)           | D|     |Qd | Max| Qd | Max | RRS Ops  |
+-----+-----+-----+-----+-----+-----+
|0x8000000500C79C00|196300640006002B|H| 0x00|0000|0072|0000|0001|      2677|
|0x8000000500C22F00|1963006400060027|H| 0x00|0000|0072|0000|0001|      1850|
|0x8000000500C14D80|196300640006002F|H| 0x00|0000|0072|0000|0001|      2657|
|0x8000000500C1AD80|1963006400060023|H| 0x00|0000|0072|0000|0001|      3546|
[...]
+-----+-----+-----+-----+
Avg|   RRS|   RRS | Large|
RRS|   TLF|   Read| Chains|
+-----+-----+-----+-----+
 39| 34819|  9791|      0|
 25| 42323|  6798|      0|
 35| 37134|  7371|      0|
 44| 31723|  8203|      0|
[...]
(Output split and truncated)
```

```
switch:admin> portshow xtun 23 -ficon -images
FCUB (FICON Control Unit Block) Count = 11

(0x)          Side Path:          devRange DevCnt Model
=====
8000000500C82600  H 0x19630064000601**      -      0x00  3990
8000000500BFAB00  H 0x19630064000600**    20-BF  0x80  3990
8000000500C81F00  D 0x19640F63040601**    00-07  0x08  3088
8000000500C83300  H 0x196302640A0510**    00-0F  0x10  3590
```

FC Egress parms:

```

Type   LPE BL CHPID          Tag V Ve P B P VC EGID cbState
====   == == =====
DISK   Yes   0000  0x02908700 Y  Y 2 0 0 05 0010  0
DISK   Yes   0000  0x0290E841 Y  Y 2 0 0 05 0010  0
FCTC   Yes   0000  0x02C4D969 Y  Y 2 0 0 15 0013  0
TAPE   Yes   0000  0x023D22B4 Y  Y 2 0 3 15 0012  0
(Output split and truncated)

```

To display FICON emulation statistics for the FCIP tunnel:

```
switch:admin> portshow xtun 23 -ficon -stats
```

```

FICON FCIP Tunnel=7 Statistics:
Emulation Common Statistical Counts (decimal)
=====
TotalIngressFrames      =          1824707656
TotalEgressFrames       =          1665499614
TotalFCEgressFrames     =          136476952
TotalCmds                =          152548501
TotalEmulDvcLvlAcks     =           81667
TotalEmulatedOps        =          23893981
Idle Accepted Status    =          8000016
Idle Not Accpt Status   =          11886003
MaxRetryQueueDepth     = 0      MaxEgressQueueDepth =          2316
TotalCUBusyResponses   = 0      TotalCUEndResponses =          0
TotalEmulatedCUBusys   = 0      TotalEmulCUEnd       =          0
TotalSelectiveResets   = 0      TotalChLinkBusy      =          0
TotalCancels           = 0      TotalAborts          =          0
TotalEmulErrors        = 0      TotalCuLinkBusy      =          0
TotalPurgePaths        = 0      Xport LRC CheckErrors=          0
Generated Link Busys   = 0      Failed Generate Frame=          0

WIRE Buffer Percentages LocalFree= 95 LocalLow= 94 PeerFree= 95 PeerLow=93
XBAR Buffer Percentages LocalFree= 98 LocalLow= 97 PeerFree= 98 PeerLow=97

```

```
FICON Configuration Status:
```

```

-----
FICON Host side Paths = 3 Device Side Paths= 1
FICON LPARS Connected = 5 Curr LCUs = 17 Total Devices= 200
Current Egress Q Count= 0 RetryQCnt = 0
Current Free Headers = 892 HdrsInUse = 0
Active Emulation Count= 4

```

```
XRC Emulation Statistical Counts (decimal)
```

```

-----
Total Emulated RRS Chains      =          1950996
Total Emulated RRS Commands    =          18732495
Total Received RRS Bytes       =          145156149548
Total XRC RRS Requests in Bytes =          315684784416
Average RRS Request BlkSize    =           7748
Average RRS Requests per Chain =           9
Largest RRS Request Byte Count =          12714240
Ratio of RRS Read Bytes to the
Requested Read bytes = 459 read : 1000 requested

```

```
Tape Write Emulation Statistical Counts (decimal)
```

```

-----
Current Bytes in write pipe      =          0
Maximum Bytes in write pipe     =          21359520
Largest write chain processed    =          4128769

```

| | |
|---------------------------------------|---------------|
| Total number of emulated Write Bytes= | 1374007085422 |
| Number of emulated Write Chains = | 11665734 |
| Total number of emulated Write CCWs = | 39204146 |
| Average Emulated Writes Blocksize = | 35047 |
| Average Writes in Emulated Chains = | 3 |
| Write emulation slowdowns = | 138931 |
| Slowdowns at Start of Chain = | 83938 |
| Slowdowns at End of chain = | 54993 |
| Current Host side Write FDCB Count = | 4 |
| Max Concurrent Write FDCB Count = | 16 |
| Current Write Limited FDCB Count = | 0 |

Tape Read Emulation Statistical Counts (decimal)

| | |
|---------------------------------------|---------------|
| Total number of emulated Read Bytes = | 1299985900584 |
| Total number of emulated Read Chains= | 10277251 |
| Total number of emulated Read CCWs = | 35565516 |
| Average Emulated Bytes per chain = | 126491 |
| Average Emulated Read Blocksize = | 36551 |
| Average CCWs in Emulated Chains = | 3 |
| Current Read FDCBs Count = | 0 |
| Max Concurrent Read FDCB Count = | 16 |
| Current Read Limited FDCB Count = | 0 |

Tera Write Emulation Statistical Counts (decimal)

| | |
|---------------------------------------|-------------|
| Current Bytes in write pipe = | 0 |
| Maximum Bytes in write pipe = | 1405952 |
| Largest write chain processed = | 43936 |
| Total number of emulated Write Bytes= | 19204226768 |
| Number of emulated Write Chains = | 989882 |
| Total number of emulated Write CCWs = | 989882 |
| Average Emulated Writes Blocksize = | 19400 |
| Average Writes in Emulated Chains = | 1 |
| Write emulation slowdowns = | 114 |
| Slowdowns at Start of Chain = | 0 |
| Slowdowns at End of chain = | 114 |
| Single Chain Emulation Counter = | 0 |
| Write Paced Count = | 0 |
| Current Host side Write FDCB Count = | 1 |
| Max Concurrent Write FDCB Count = | 1 |
| Current Write Limited FDCB Count = | 0 |

Tera Read Emulation Statistical Counts (decimal)

| | |
|---------------------------------------|-------------|
| Total number of emulated Read Bytes = | 17688908128 |
| Total number of emulated Read Chains= | 825572 |
| Total number of emulated Read CCWs = | 825572 |
| Average Emulated Bytes per chain = | 21426 |
| Average Emulated Read Blocksize = | 21426 |
| Average CCWs in Emulated Chains = | 1 |
| Read Block Paced Count = | 0 |
| Read Not Ready Situations Count = | 0 |
| Current Read FDCBs Count = | 1 |
| Max Concurrent Read FDCB Count = | 1 |
| Current Read Limited FDCB Count = | 1 |

Printer Write Emulation Statistical Counts (decimal)

| | |
|-------------------------------|---|
| Current Bytes in write pipe = | 0 |
|-------------------------------|---|

```

Maximum Bytes in write pipe      =          45627
Largest write chain processed    =           9150
Total number of emulated Write Bytes=      27375500
Number of emulated Write Chains =          3500
Total number of emulated Write CCWs =          6500
Average Emulated Writes Blocksize =          4211
Average Writes in Emulated Chains =           1
Current Host side Write FDCB Count =           1
Max Concurrent Write FDCB Count  =           1
Current Write Limited FDCB Count  =           1

```

FICON Debug Flags (ftrace and others) = 0xffc98030 (Default = 0xf7c90000)

```

-----
Bit 31 [0x80000000] = 1 : TRIGGER_ON_SELRESET
Bit 30 [0x40000000] = 1 : TRIGGER_ON_PURGEPATH
Bit 29 [0x20000000] = 1 : TRIGGER_ON_RRS_MISS
Bit 28 [0x10000000] = 1 : TRIGGER_ON_LRJ

Bit 27 [0x08000000] = 1 : TRIGGER_ON_UNIT_CHECK
Bit 26 [0x04000000] = 1 : TRIGGER_ON_LOOKUP_FAIL
Bit 25 [0x02000000] = 1 : TRIGGER_ON_FDCB_ABORT
Bit 24 [0x01000000] = 1 : TRIGGER_ON_NOFDCB_ABORT

Bit 23 [0x00800000] = 1 : TRIGGER_ON_LINKDOWN
Bit 22 [0x00400000] = 1 : TRIGGER_ON_SENSE_CMD
Bit 21 [0x00200000] = 0 : TRIGGER_ON_BUSYATTN
Bit 20 [0x00100000] = 0 : TRIGGER_ON_XRCUNS

Bit 19 [0x00080000] = 1 : DISPLAY_FDCB_ON_ERROR
Bit 18 [0x00040000] = 0 : TRIGGER_ON_LOW_WIREPOOL
Bit 17 [0x00020000] = 0 : TRIGGER_ON_READ_UNITCHECK
Bit 16 [0x00010000] = 1 : TRIGGER_ON_STATE_SAVE

Bit 15 [0x00008000] = 1 : TRIGGER_ON_SUSPEND
Bit 11 [0x00000800] = 0 : FORCE_DISABLE_3203_PRINT
Bit 10 [0x00000400] = 0 : FORCE_DISABLE_3211_PRINT
Bit 9  [0x00000200] = 0 : FORCE_DISABLE_3800_PRINT
Bit 8  [0x00000100] = 0 : FORCE_DISABLE_3900_PRINT

Bit 4  [0x00000010] = 1 : ACCEPT_READTRACK_STATUS

Bit 3  [0x00000008] = 0 : FORCE_DISABLE_READ_PIPE
Bit 2  [0x00000004] = 0 : FORCE_DISABLE_WRITE_PIPE
Bit 1  [0x00000002] = 0 : FORCE_DISABLE_XRC_EMUL
Bit 0  [0x00000001] = 0 : DISABLE_PERSIST_IU_PACE

```

To display FICON Teradata Read and Write performance data:

```
switch:admin> portshow xtun 16 -ficon -teraperf
```

```

Tera Performance Monitor Data:
=====
      Sample Time Period in ms: 4318

Tera Write Performance Data:
-----
      Emulated Chains per sec: 1186
      Emulated CCWs per sec: 1186
      Ave Write Block Size: 22617
      Emulated Write BPS: 26831514

```

portShow

```
Tera Read Performance Data:
-----
Emulated Chains per sec: 1069
Emulated CCWs per sec: 1069
Ave Read Block Size: 10780
Emulated Read BPS: 11531104
```

To display FICON Printer performance data:

```
switch:admin> portshow xtun 23 -ficon -printperf
```

```
Printer Performance Monitor Data:
=====
Sample Time Period in ms: 7460
-----
No Emulated Printer operations
-----
No Emulated Read operations
```

To display FICON tape performance data:

```
switch:admin> portshow xtun 23 -ficon -tapeperf
```

```
Tape Performance Monitor Data:
=====
Sample Time Period in ms: 2461
Tape Write Performance Data:
-----
Emulated Chains per sec: 480
Emulated CCWs per sec: 3069
Ave Write Block Size: 32760
Emulated Write BPS: 100540440
-----
No Emulated Read operations
```

To display FICON XRC performance data:

```
switch:admin> portshow xtun 23 -ficon -xrcperf
```

```
XRC Performance Monitor Data:
=====
Sample Time Period in ms: 11340
Emulated Chains per sec: 16
Emulated RRS Cmds per sec: 16
Emulated RRS Bytes per sec: 640
Average RRS Update Size: 40
```

SEE ALSO [authUtil](#), [portCfg](#), [portLoginShow](#), [portCfgLongDistance](#), [switchShow](#)

portStats64Show

Displays the 64-bit hardware statistics for a port.

Synopsis `portstats64show [slot]port [-long]`

Description Use this command to display 64-bit hardware statistics for a specified port. When used without the **-long** option, two integers are reported for most values, the lower and upper 32-bits are reported as two separate numbers. In this case, the top word is the most significant. When issued with the **-long** option, the command displays the counters as one single 64-bit number.

| | |
|----------------------|--|
| stat64_wtx | Number of 4-byte words transmitted. |
| stat64_wrx | Number of 4-byte words received. |
| stat64_ftx | Number of frames transmitted. |
| stat64_frx | Number of frames received. |
| stat64_c2_frx | Number of class 2 frames received. |
| stat64_c3_frx | Number of class 3 frames received. |
| stat64_lc_rx | Number of link control frames received. |
| stat64_mc_rx | Number of multicast frames received. |
| stat64_mc_to | Number of multicast timeouts. |
| stat64_mc_tx | Number of multicast frames transmitted. |
| tim64_rdy_pri | Number of times R_RDY was high priority. |
| tim64_txcrd_z | Number of times that the TX BB_credit was at zero. |
| er64_enc_in | Number of encoding errors inside of frames. |
| er64_crc | Number of frames with CRC errors. |
| er64_trunc | Number of frames shorter than minimum. |

| | |
|----------------------------------|--|
| er64_toolong | Number of frames longer than maximum. |
| er_bad_eof | Number of frames with bad end-of-frame. |
| er64_enc_out | Number of encoding error outside of frames. |
| er64_disc_c3 | Number of class 3 frames discarded. |
| er64_pcs_blk | Number of Physical Coding Sublayer (PCS) block errors. This counter records encoding violations on 10 Gbps or 16 Gbps ports. |
| stat64_fec_cor | The number of errors corrected by FEC. Displayed only on 16 Gbps-capable platforms and not on 32 Gbps-capable platforms. |
| stat64_fec_uncor | The number of errors left uncorrected by FEC. |
| stat64_rateTxFrame | Tx frame rate (frames/second). |
| stat64_rateRxFrame | Rx frame rate (frames/second). |
| stat64_rateTxPeakFrame | Tx peak frame rate (frames/second). |
| stat64_rateRxPeakFrame | Rx peak frame rate (frames/seconds). |
| stat64_rateTxWord | Tx Word rate (words/seconds). |
| stat64_rateRxWord | Rx Word rate (words/seconds). |
| stat64_rateTxPeakWord | Tx peak Word rate (words/sec). |
| stat64_rateRxPeakWord | Rx peak Word rate (words/sec). |
| stat64_aveTxFrameSize> | Average Tx Frame size |
| stat64_aveRxFrameSize | Average Rx Frame size |
| stat64_PRJTFramees | Number of P_RJT frames transmitted. |
| stat64_PBSYFrames | Number of P_BSY transmitted. |

stat64_inputBuffersFull

Number of occasions on which input buffers are full.

stat64_rxClass1Frames

Number of class 1 frames received.

stat64_aveTxFrameSize

Average Tx Frame size, based on the word and frame counts during the last five seconds.

stat64_aveRxFrameSize

Average Rx Frame size, based on the word and frame counts during the last five seconds.

The following counters provided by SNMP are displayed with **portStats64Show -long** on switches running Fabric OS v6.4.0 or later:

swConnUnitZeroTenancy, zero_tenancy

Number of times a zero tenancy occurred.

swConnUnitFLNumOfTenancy, fl_tenancy

Number of times the FL_Port had a loop tenancy.

swConnUnitNLNumOfTenancy, nl_tenancy

Number of times any NL_Port had a loop tenancy.

swConnUnitStopTenancyStarvation, Starve_stop

Number of loop tenancies stopped due to starvation.

swConnUnitOpen, opened

Number of times the FL_Port entered OPENED state.

swConnUnitTransferConnection, transfer

Number of times the FL_Port entered TRANSFER state.

swConnUnitOpen, open

Number of times the FL_Port entered OPEN state.

swConnUnitInvalidARB, er_inv_arb

Number of invalid arbitrated loops (ARBs).

The following counters are platform-specific and applicable only to 8 Gbps-capable ASICs only. Refer to the *Fabric OS Troubleshooting and Diagnostics Guide*, Appendix A, for a table that correlates ASIC type with switch models.

swConnUnitFTB1Miss, er_type1_miss

The number of FCR frames with transmit errors.

swConnUnitFTB2Miss, er_type2_miss

The number of frames with routing errors.

swConnUnitFTB6Miss, er_type6_miss

The number of FCR frames with receive errors.

swConnUnitZoneMiss, er_zone_miss

Number of frames with hard zoning miss

swConnUnitLunZoneMiss, er_lun_zone_miss

Number of frames with logical unit number (LUN) zoning miss.

swConnUnittStatRxMulticastToObjects, stat_mc_to

The number of multicast timeouts.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports.

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the port to be displayed, followed by a slash (/).

port

Specify a port number to be displayed, relative to its slot for bladed systems. Use **switchShow** to display a list of valid ports.

-long

Displays the counters as one single 64-bit number. This operand is optional; if omitted, the lower and upper 32-bits are reported as two separate numbers.

Examples To display the 64-bit hardware statistics for a port:

```
switch:user> portstats64show 4/15
```

```
stat64_wtx    0    top_int : 4-byte words transmitted
              21618 bottom_int : 4-byte words transmitted
stat64_wrx    0    top_int : 4-byte words received
              22492 bottom_int : 4-byte words received
stat64_ftx    0    top_int : Frames transmitted
              596   bottom_int : Frames transmitted
stat64_frx    0    top_int : Frames received
              614   bottom_int : Frames received
stat64_c2_frx 0    top_int : Class 2 frames received
              0     bottom_int : Class 2 frames received
stat64_c3_frx 0    top_int : Class 3 frames received
              0     bottom_int : Class 3 frames received
stat64_lc_rx  0    top_int : Link control frames received
              462   bottom_int : Link control frames received
stat64_mc_rx  0    top_int : Multicast frames received
              0     bottom_int : Multicast frames received
stat64_mc_to  0    top_int : Multicast timeouts
              0     bottom_int : Multicast timeouts
stat64_mc_tx  0    top_int : Multicast frames transmitted
              0     bottom_int : Multicast frames transmitted
tim64_rdy_pri 0    top_int : Time R_RDY high priority
              0     bottom_int : Time R_RDY high priority
tim64_txcrd_z 0    top_int : Time BB_credit zero
              34211 bottom_int : Time BB_credit zero
er64_enc_in   0    top_int : Encoding errors inside of frames
              0     bottom_int : Encoding errors inside of frames
er64_crc      0    top_int : Frames with CRC errors
              0     bottom_int : Frames with CRC errors
er64_trunc    0    top_int : Frames shorter than minimum
              0     bottom_int : Frames shorter than minimum
```

```

er64_toolong 0 top_int : Frames longer than maximum
              0 bottom_int : Frames longer than maximum
er64_bad_eof 0 top_int : Frames with bad end-of-frame
              0 bottom_int : Frames with bad end-of-frame
er64_enc_out 0 top_int : Encoding error outside of frames
              135762 bottom_int : Encoding error outside of frames
er64_disc_c3 0 top_int : Class 3 frames discarded
              0 bottom_int : Class 3 frames discarded
er64_pcs_blk 0 top_int : PCS block errors
              0 bottom_int : PCS block errors
stat64_fec_cor 0 top_int : FEC corrected errors detected
              0 bottom_int : FEC corrected errors detected
stat64_fec_uncor 0 top_int : FEC uncorrected errors detected
              0 bottom_int : FEC uncorrected errors detected
stat64_rateTxFrame 0 Tx frame rate (fr/sec)
stat64_rateRxFrame 0 Rx frame rate (fr/sec)
stat64_rateTxPeakFrame 9 Tx peak frame rate (fr/sec)
stat64_rateRxPeakFrame 9 Rx peak frame rate (fr/sec)
stat64_rateTxWord 0 Tx Word rate (words/sec)
stat64_rateRxWord 0 Rx Word rate (words/sec)
stat64_rateTxPeakWord 264 Tx peak Word rate (words/sec)
stat64_rateRxPeakWord 272 Rx peak Word rate (words/sec)
stat64_aveTxFrameSize 504 Average Tx Frame size
stat64_aveRxFrameSize 511 Average Rx Frame size
stat64_PRJTFFrames 0 top_int : 4-byte words transmitted
                  0 bottom_int : 4-byte words transmitted
stat64_PBSYFFrames 0 top_int : 4-byte words transmitted
                  0 bottom_int : 4-byte words transmitted
stat64_inputBuffersFull 0 top_int : 4-byte words transmitted
                        0 bottom_int : 4-byte words transmitted
stat64_rxClass1Frames 0 top_int : 4-byte words transmitted
                      0 bottom_int : 4-byte words transmitted
stat64_aveTxFrameSize 0 Average Tx Frame size
stat64_aveRxFrameSize 0 Average Rx Frame size

```

To display the counters as one single 64-bit number:

```

switch:admin> portstats64show 12 -long
zero64_tenancy 0 zero_tenancy
fl64_tenancy 0 number of times FL has the tenancy
nl64_tenancy 0 number of times NL has the tenancy
starve64_stop 0 tenancies stopped due to starvation
opened64 0 FL_Port opened
transfer64 0 loop_transfer
open64 0 loop_open
er64_inv_arb 0 Invalid ARB
er64_type1_miss 0 frames with FTB type 1 miss
er64_type2_miss 0 frames with FTB type 2 miss
er64_type6_miss 0 frames with FTB type 6 miss
er64_zone_miss 0 frames with hard zoning miss
er64_lun_zone_miss 0 frames with LUN zoning miss
lli64 0 Low level interrupts

```

SEE ALSO [portStatsClear](#), [portStatsShow](#)

portStatsClear

Clears port hardware statistics.

Synopsis `portstatsclear [slot/]port`
`portstatsclear -i [index1[-index2] [...] [-f]]`
`portstatsclear -x [hex1[-hex2] [...]]`
`portstatsclear -slot [slot1[-slot2][...]]`
`portstatsclear -h`

Description Use this command to clear the hardware statistics for the specified ports. Including ALPA-based CRC monitor, End-to-End monitor, and Filter-based performance monitor statistics.

You can identify a single port to be cleared by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers(decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index argument to display the **portSwap** status, or alternately use **portSwapShow**.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports.

Operands This command has the following operands:

| | |
|---------------------------|--|
| <i>slot</i> | For bladed systems only, specify the slot number of the port to be cleared, followed by a slash (/). |
| <i>port</i> | Clears a single port identified by the port number, relative to its slot on bladed systems. Port ranges are not supported with this command. Use switchShow for a listing of valid ports. |
| <i>-i index1[-index2]</i> | Clears a single port or a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, -i 33-47 65-73. |
| <i>-f</i> | Ignores nonexistent ports. This operand is valid only with the -i option. |
| <i>-x [hex1 [-hex2]]</i> | Clears a port or a range of ports identified by port numbers, index number in hexadecimal format. You may specify multiple port ranges separated by a space, for example, -x 21-26 28-3c. |

-slot [*slot1*[-*slot2*]

Clears all ports on a slot or on a range of slots, for example, **-s 3-5**. You may specify multiple slot ranges separated by a space, for example, **-s 3-5 8-10**.

-h

Displays the command usage.

Examples

To clear hardware statistics for a single port specified by its port number:

```
switch:admin> portstatsclear 4/15
```

To clear hardware statistics for a single port specified by its index number:

```
switch:admin> portstatsclear -i 25
```

To clear hardware statistics for a range of ports specified by their index numbers:

```
switch:admin> portstatsclear -i 32-40
```

To clear hardware statistics for multiple port ranges specified by their index numbers:

```
switch:admin> portstatsclear -i 32-40 50-56
```

To clear hardware statistics for all ports on slots 3-5:

```
switch:admin> portstatsclear -s 3-5
```

To clear hardware statistics for all ports on slots 3-5 and 7-10:

```
switch:admin> portstatsclear -s 3-5 7-10
```

To clear hardware statistics for range of ports specified in hexadecimal format:

```
switch:admin> portstatsclear -x 1d-1e
```

SEE ALSO [portStats64Show](#), [portStatsShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portStatsShow

Displays port hardware statistics.

Synopsis

```
portstatsshow [slot]port
portstatsshow -i [index1[-index2]][...] [-f]]
portstatsshow -x [hex1[-hex2] [...]]
portstatsshow -slot [slot1[-slot2]][...]
portstatsshow ge [slot]ge port
portstatsshow ip [slot]ge port [ip_address]
portstatsshow fcip [slot]ge port [tunnel_number]
portstatsshow -h
```

Description Use this command to display port hardware statistics counters. Some counters are platform- or port-specific and display only on those platforms and ports. All statistics have a maximum 32-bit value of 4,294,967,295 except stat_wtx and stat_wrx. The stat_wtx and stat_wrx fields have a maximum 64-bit value of 18,446,744,073,709,551,615.

You can display statistics of a single port to be cleared by its port number or by its port index number in decimal or hexadecimal format. Port ranges are supported with port numbers, index numbers(decimal or hexadecimal) or by specifying a slot or a slot range. Use **switchShow** for a listing of valid ports, slots, and port index numbers.

Specifying multiple ports with the index (-i) or slot (-s) option is supported only if **PortSwap** is disabled. They are not supported on GbE ports and configured F_Port trunks. Use the -i option without a port index to display the **portSwap** status, or alternately use **portSwapShow**.

The command output may include the following fields (Tx indicates frames transmitted by the port; Rx indicates frames received by the port).

| | |
|--------------------|---|
| stat_wtx | The number of 4-byte words transmitted. |
| stat_wrx | The number of 4-byte words received. |
| stat_ftx | The number of class 2, class 3, and control frames transmitted. |
| stat_frx | The number of class 2, class 3, and control frames received. |
| stat_c2_frx | The number of class 2 frames received. |
| stat_c3_frx | The number of class 3 frames received. |

| | |
|-------------------------|---|
| stat_lc_rx | The number of link control frames received. |
| stat_mc_rx | The number of multicast frames received. |
| stat_mc_to | The number of multicast timeouts. |
| stat_mc_tx | The number of multicast frames transmitted. |
| tim_rdy_pri | The number of times that sending R_RDY or VC_RDY primitive signals was a higher priority than sending frames, due to diminishing credit reserves in the transmitter at the other end of the fiber. This parameter is sampled at intervals of 1.8 microseconds, and the counter is incremented by 1 if the condition is true. |
| tim_txcrd_z | The number of times that the port was unable to transmit frames because the transmit BB credit was zero. The purpose of this statistic is to detect congestion or a device affected by latency. This parameter is sampled at intervals of 2.5 microseconds, and the counter is incremented if the condition is true. Each sample represents 2.5 microseconds of time with zero Tx BB Credit. An increment of this counter means that the frames could not be sent to the attached device for 2.5 microseconds, indicating degraded performance. |
| tim_txcrd_z_vc | The number of times that the port was unable to transmit frames because the transmit BB credit was zero for each of the port's 16 Virtual Channels (VC 0-15). The purpose of this statistic is to detect congestion or a device affected by latency. This parameter is sampled at intervals of 2.5 microseconds (microseconds), and the counter is incremented if the condition is true. Each sample represents 2.5 microseconds of time with zero Tx BB Credit. An increment of this counter means that the frames could not be send to the attached device for 2.5 microseconds, indicating degraded performance (platform- and port-specific). |
| tim_latency_vc | TxQ latency monitor statistics. Displayed only on 16 Gbps-capable platforms and not on 32 Gbps-capable platforms. |
| lat_tot_pkt_vc | Total TxQ latency monitor statistics. Displayed only on 32 Gbps-capable platforms. |
| lat_hi_time_vc | High TxQ latency monitor statistics. Displayed only on 32 Gbps-capable platforms. |
| lat_lo_time_vc | Low TxQ latency monitor statistics. Displayed only on 32 Gbps-capable platforms. |
| latency_dma_ts | TXQ latency DMA timeStamp. Displayed only on 32 Gbps-capable platforms. |
| fec_cor_detected | The number of errors corrected by FEC. Displayed only on 16 Gbps-capable platforms and not on 32 Gbps-capable platforms. |

| | |
|---------------------------|--|
| fec_uncor_detected | The number of errors left uncorrected by FEC. |
| er_enc_in | The number of encoding errors inside frames. |
| er_crc | The number of frames with cyclic redundancy check (CRC) errors. |
| er_trunc | The number of frames shorter than the minimum frame length. |
| er_toolong | The number of frames longer than the maximum frame length. |
| er_bad_eof | The number of frames with bad end-of-frame. |
| er_enc_out | The number of encoding error outside frames. |
| er_bad_os | The number of invalid ordered sets (platform- and port-specific). |
| er_pcs_blk | The number of Physical Coding Sublayer (PCS) block errors. This counter records encoding violations on 10 Gbps or 16 Gbps ports. This is applicable only on platforms that support 10 Gbps or 16 Gbps ports. |
| er_rx_c3_timeout | The number of receive class 3 frames received at this port and discarded at the transmission port due to timeout (platform-and port-specific). |
| er_tx_c3_timeout | The number of transmit class 3 frames discarded at the transmission port due to timeout (platform- and port-specific). |
| er_unroutable | The number of frames discarded because they cannot be routed. |
| er_unreachable | The number of frames discarded because the destination port cannot be reached. This field is specific to 16 Gbps-capable platforms. |
| er_c3_dest_unreach | The number of class 3 frames discarded because the destination cannot be reached. This field is specific to 8 Gbps-capable platforms. |
| er_other_discard | The number of other discarded due to route lookup failures or other reasons. |
| er_zone_discard | The number of class 3 frames discarded due to zone mismatch. |
| er_type1_miss | The number of FCR frames with transmit errors. |

| | |
|--|---|
| er_type2_miss | The number of frames with routing errors. |
| er_type6_miss | The number of FCR frames with receive errors. |
| er_zone_miss , er_lun_zone_miss | The number of frames discarded due to hard zoning miss or LUN zoning miss. If Rx port hard zoning is enabled, frames will be discarded at the Rx port. If TX port hard zoning is enabled, frames will be discarded at the TX port. If both RX and TX port hard zoning is enabled, frames will be discarded at the RX port. (LUN zoning is currently not supported.) |
| er_crc_good_eof | The number of CRC errors with good end-of-frame (EOF) (platform- and port-specific). |
| er_inv_arb | The number of invalid arbitrated loops (ARBs). |
| er_single_credit_loss | The number of times the port lost a single VC_RDY primitive signal or a single frame. |
| er_multi_credit_loss | The number of times the port lost multiple VC_RDY primitive signals or multiple frames. |
| fec_corrected_rate | The number of FEC corrected blocks per second. Displayed only on 32 Gbps-capable platforms. |
| compression_ratio | The ratio based on the last 5 second sampled accumulated value. It is represented in percentage. If the port is configured for compression, this value is displayed. |
| open | The number of times the FL_Port entered OPEN state. |
| transfer | The number of times the FL_Port entered TRANSFER state. |
| opened | The number of times the FL_Port entered OPENED state. |
| starve_stop | The number of loop tenancies stopped due to starvation. |
| fl_tenancy | The number of times the FL_Port had a loop tenancy. |
| nl_tenancy | The number of times the NL_Port had a loop tenancy. |
| zero_tenancy | The number of times a zero tenancy occurred. |
| ge_stat_tx_frms | The number of frames transmitted on the GbE port. |
| ge_stat_tx_octets | The number of octets transmitted on the GbE port. |

| | |
|------------------------------|--|
| ge_stat_tx_ucast_frms | The number of unicast frames transmitted on the GbE port. |
| ge_stat_tx_mcast_frms | The number of multicast frames transmitted on the GbE port. |
| ge_stat_tx_bcast_frms | The number of broadcast frames transmitted on the GbE port. |
| ge_stat_tx_vlan_frms | The number of VLAN frames transmitted on the GbE port. |
| ge_stat_tx_pause_frms | The number of pause frames transmitted on the GbE port. |
| ge_stat_rx_frms | The number of frames received on the GbE port. |
| ge_stat_rx_octets | The number of octets received on the GbE port. |
| ge_stat_rx_ucast_frms | The number of unicast frames received on the GbE port. |
| ge_stat_rx_mcast_frms | The number of multicast frames received on the GbE port. |
| ge_stat_rx_bcast_frms | The number of broadcast frames received on the GbE port. |
| ge_stat_rx_vlan_frms | The number of VLAN frames received on the GbE port. |
| ge_stat_rx_pause_frms | The number of pause frames received on the GbE port. |
| ge_err_carrier | The number of times the GbE port lost carrier sense. |
| ge_err_length | The number of times an invalid length error was observed on the GbE port. |
| ge_err_crc | The number of CRC Errors received on the GbE port. |
| ge_err_abort | The number of frames aborted on the GbE port. |
| ge_err_overnun | The number of overruns observed on the GbE port. |
| ge_err_fifo_ovf | The number of times an overflow of the first in first out (FIFO) queue was observed on the GbE port. |
| ip_err_hdr_cksum | The number of checksum errors observed on the GbE port. |

ip_err_tcp_data_chksum

The number of IP TCP data checksum errors observed on the GbE port.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is not supported on FCoE ports.

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to be displayed, followed by a slash (/).

[ge]port

Displays statistics for a single port identified by the port number, relative to its slot on bladed systems. Specify the optional **ge** option to display the GbE port hardware statistics. Port ranges are not supported with this command. Use **switchShow** for a listing of valid ports.

-i index1[-index2]

Displays statistics for a single port or for a range of ports identified by port index numbers. You may specify multiple index ranges separated by a space, for example, **-i 33-47 65-73**.

-f

Ignores nonexisting ports. This operand is valid only with the **-i** option.

-x [hex1 [-hex2]]

Specifies a port or a range of ports identified by port index numbers in hexadecimal format. You may specify multiple port ranges separated by a space, for example, **-x 21-26 28-3c**.

-slot [slot1[-slot2]]

Displays statistics for all ports on a slot or on a range of slots, for example, **-s 3-5**. You may specify multiple slot ranges separated by a space, for example, **-s 3-5 8-10**.

ge

Displays the GbE port statistics.

ip

Displays all GbE port statistics related to IP addresses that are not zero. This operand is not supported on the Brocade 7800 and FX8-24 platforms.

ip_address

Specifies an IP address to display statistics only for the specified IP address. This operand is optional and valid only with the **ip** option.

fcip

Displays the GbE statistics on all FCIP tunnels. This operand is not supported on the Brocade 7800 and FX8-24 platforms.

tunnel_number

Specifies a tunnel ID to display statistics only for the specified FCIP tunnel. This operand is optional and valid only with the **fcip** option.

-h

Displays the command usage.

Examples To display the basic set of statistics for port 1/13 on a 16 Gbps-capable platform:

```

switch:admin> portstatsshow 1/13
stat_wtx          0          4-byte words transmitted
stat_wrx          0          4-byte words received
stat_ftx          0          Frames transmitted
stat_frx          0          Frames received
stat_c2_frx       0          Class 2 frames received
stat_c3_frx       0          Class 3 frames received
stat_lc_rx        0          Link control frames received
stat_mc_rx        0          Multicast frames received
stat_mc_to        0          Multicast timeouts
stat_mc_tx        0          Multicast frames transmitted
tim_rdy_pri       0          Time R_RDY high priority
tim_txcrd_z       0          Time TX Credit Zero (2.5Us ticks)
tim_txcrd_z_vc   0- 3: 0          0          0          0
tim_txcrd_z_vc   4- 7: 0          0          0          0
tim_txcrd_z_vc   8-11: 0         0          0          0
tim_txcrd_z_vc  12-15: 0         0          0          0
tim_latency_vc   0- 3: 1          1          1          1
tim_latency_vc   4- 7: 1          1          1          1
tim_latency_vc   8-11: 1         1          1          1
tim_latency_vc  12-15: 1         1          1          1
lat_tot_pkt_vc   0- 3: 1          1          1          1
lat_tot_pkt_vc   4- 7: 1          1          1          1
lat_tot_pkt_vc   8-11: 1         1          1          1
lat_tot_pkt_vc  12-15: 1         1          1          1
lat_hi_time_vc   0- 3: 0          0          0          0
lat_hi_time_vc   4- 7: 0          0          0          0
lat_hi_time_vc   8-11: 0         0          0          0
lat_hi_time_vc  12-15: 0         0          0          0
lat_lo_time_vc   0- 3: 1          1          1          1
lat_lo_time_vc   4- 7: 1          1          1          1
lat_lo_time_vc   8-11: 1         1          1          1
lat_lo_time_vc  12-15: 1         1          1          1
max_latency_vc   0- 3: 1          1          1          1
max_latency_vc   4- 7: 1          1          1          1
max_latency_vc   8-11: 1         1          1          1
max_latency_vc  12-15: 1         1          1          1
latency_dma_ts   05-11-2015 UTC Thu 02:19:16      TXQ Latency DMA TimeStamp
fec_cor_detected  0          Count of blocks that were corrected by FEC
fec_uncor_detected 0          Count of blocks that were left uncorrected by FEC
er_enc_in        0          Encoding errors inside of frames
er_crc           0          Frames with CRC errors
er_trunc         0          Frames shorter than minimum
er_toolong       0          Frames longer than maximum
er_bad_eof       0          Frames with bad end-of-frame
er_enc_out       0          Encoding error outside of frames
er_bad_os        0          Invalid ordered set
er_rx_c3_timeout 0          Class 3 receive frames discarded due to timeout
er_tx_c3_timeout 0          Class 3 transmit frames discarded due to timeout
er_c3_dest_unreach 0          Class 3 frames discarded due to destination
  unreachable
er_other_discard 0          Other discards
er_type1_miss    0          frames with FTB type 1 miss
er_type2_miss    0          frames with FTB type 2 miss
er_type6_miss    0          frames with FTB type 6 miss
er_zone_miss     0          frames with hard zoning miss
er_lun_zone_miss 0          frames with LUN zoning miss
er_crc_good_eof  0          Crc error with good eof

```

```

er_inv_arb          0          Invalid ARB
compression_ratio   6          Average compression ratio
open                0          loop_open
transfer            0          loop_transfer
opened              0          FL_Port opened
starve_stop         0          tenancies stopped due to starvation
fl_tenancy          0          number of times FL has the tenancy
nl_tenancy          0          number of times NL has the tenancy
zero_tenancy        0          zero tenancy
phy_stats_clear_ts  04-22-2013 MDT Mon 17:08:41  Timestamp of phy_port stats clear
lgc_stats_clear_ts  04-22-2013 MDT Mon 17:08:41  Timestamp of lgc_port stats clear

```

To display the basic set of statistics using port index numbers:

```

switch:admin> portstatsshow -i 13
switch:admin> portstatsshow -i 13-23
switch:admin> portstatsshow -i 4-6 22-30

```

To display the basic set of statistics using slot numbers:

```

switch:admin> portstatsshow -s 3-5
switch:admin> portstatsshow -s 3-5 10-13

```

To display GbE port statistics for GbE1 on the Brocade 7800:

```

switch:admin> portstatsshow ge 8/ge1
ge_stat_tx_frms      1523916  GE transmitted frames
ge_stat_tx_octets     152411630  GE transmitted octets
ge_stat_tx_ucast_frms 1523907   GE transmitted unicast frames
ge_stat_tx_mcast_frms 0         GE transmitted multicast frames
ge_stat_tx_bcast_frms 9         GE transmitted broadcast frames
ge_stat_tx_vlan_frms 0         GE transmitted vlan frames
ge_stat_tx_pause_frms 0         GE transmitted pause frames
ge_stat_rx_frms       1512154  GE received frames
ge_stat_rx_octets     149255230  GE received octets
ge_stat_rx_ucast_frms 1512154  GE received unicast frames
ge_stat_rx_mcast_frms 0         GE received multicast frames
ge_stat_rx_bcast_frms 0         GE received broadcast frames
ge_stat_rx_vlan_frms 0         GE received vlan frames
ge_stat_rx_pause_frms 0         GE received pause frames
ge_err_carrier        0         GE lost carrier sense
ge_err_length         0         GE invalid length
ge_err_crc             0         GE CRC Errors
ge_err_abort           0         GE abort frames
ge_err_overrun         0         GE overruns
ge_err_fifo_ovf        0         GE Fifo overflow

```

To display the basic set of statistics using port index number specified in hexadecimal format:

```

switch:admin> portstatsshow -x 11f
port: 287
=====
stat_wtx              422          4-byte words transmitted
stat_wrx              560          4-byte words received
stat_ftx              16           Frames transmitted
stat_frx              16           Frames received
(output is truncated)

```

SEE ALSO [portErrShow](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#)

portSwap

Swaps two ports or removes swapping of ports.

Synopsis `portswap [slot1/]port1 [slot2/]port2`

`portswap --restore`

Description Use this command to swap the 24-bit port address (PID) for a pair of ports or to remove swapping of ports. When swapping both ports must be disabled prior to executing this command and the port-swapping feature must be enabled using **portSwapEnable**.

The result of this operation is persistent across reboots and power cycles. Use **portSwapShow** to display the swapped ports along with their new port address

Port swap information is kept in its own database; it cannot be manipulated by editing the configuration database displayed by **configShow** and **configUpload**. To undo a previous port swap, execute **portSwap** again on the same two ports.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

TI zoning is not supported on E_Ports or F_Ports that are configured as part of a TI Zone with ports swapped. To work around this issue, reconfigure TI zoning rather than swapping the ports.

This command is not allowed if one or both the ports are from a faulty blade.

Operands This command has the following operands:

slot1

For bladed systems only, specifies the slot number of the first port whose area number is to be swapped, followed by a slash (/).

port1

Specifies the number of the first port whose address is to be swapped, relative to its slot for bladed systems. Use **switchShow** to display a listing of valid ports.

slot2

For bladed systems only, specifies the slot number of the second port whose area number is to be swapped, followed by a slash (/).

port2

Specifies the number of the second port whose port address is to be swapped, relative to its slot for bladed systems.

--restore

Removes swapping of all ports.

Examples To swap area numbers between a pair of ports:

```
switch:admin> portswap 8/1 8/2
portswap done
```

```
switch:admin> portswapshow
PortSwap is enabled
Slot   Slotport      Swport      Address
=====
8      1              193         0x01c200
8      2              194         0x01c100
```

To cancel the swapping of ports:

```
switch:admin> portswap --restore
portswap done
switch:admin>
```

SEE ALSO [portDisable](#), [portEnable](#), [portShow](#), [portSwapDisable](#), [portSwapEnable](#), [portSwapShow](#), [switchShow](#)

portSwapDisable

Disables the PortSwap feature.

Synopsis `portswapdisable`

Description Use this command to disable the PortSwap feature. The **portSwap** command cannot be used after this feature is disabled.

The disabled state of the PortSwap feature is persistent across reboots and power cycles.

Enabling or disabling the PortSwap feature does not affect previously performed PortSwap operations.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the PortSwap feature:

```
switch:admin> portswapdisable
```

SEE ALSO [portSwap](#), [portDisable](#), [portEnable](#), [portShow](#), [portSwapEnable](#), [portSwapShow](#), [switchShow](#)

portSwapEnable

Enables the PortSwap feature.

Synopsis `portswapenable`

Description Use this command to enable the PortSwap feature. The **portSwap** command cannot be used unless the feature is first enabled with this command.

The enabled state of the PortSwap feature is persistent across reboots and power cycles.

Enabling or disabling the PortSwap feature does not affect previously performed PortSwap operations.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable the PortSwap feature:

```
switch:admin> portswapenable
```

SEE ALSO [portSwap](#), [portDisable](#), [portEnable](#), [portShow](#), [portSwapDisable](#), [portSwapShow](#), [switchShow](#)

portSwapShow

Displays the state of the PortSwap feature.

Synopsis `portswapshow`

Description Use this command to display the state of the PortSwap feature and information about swapped ports. If **portSwap** is enabled and ports have been swapped, the command displays the enabled status as well as the swapped ports and the new port address (the 24-bit PID) for these ports. If **portSwap** is disabled, the command shows the disabled status and indicates whether or not swapped ports exist on the switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display a switch with PortSwap enabled and information for swapped ports:

Console message:

```
2009/10/09-21:22:45, [PSWP-1001], 482, SLOT 6 | FID 51, INFO,
DCX_105_51, PID for port 12/0 and port 12/63 are swapped.
New PID for port 12/0 is 0x697000 and port 12/63 is 0x69e3c0.
```

```
switch:admin> portswapshow
```

```
PortSwap is enabled
Slot   Slotport      Swport      Address
=====
12     0              112         0x697000
12     63             895         0x69e3c0
```

To display the portSwap status on a switch when the feature is disabled and no ports are swapped:

```
switch:admin> portswapshow
PortSwap is disabled.
Existing Portswap condition is still effective.
Only future Portswap operations are not allowed.
```

```
No ports have been swapped
```

SEE ALSO [portSwap](#), [portDisable](#), [portEnable](#), [portShow](#), [portSwap](#), [portSwapDisable](#), [portSwapEnable](#), [switchShow](#)

portTest

Performs a functional test of a switch in a live fabric.

Synopsis

```
porttest
  [-ports itemlist]
  [-iteration count]
  [-userdelay time]
  [-timeout time]
  [-pattern pattern]
  [-patsize size]
  [-seed seed]
  [-listtype porttype]
```

Description

Use this command to isolate problems in a single replaceable element and to trace problems to near-end terminal equipment, far-end terminal equipment, or the transmission line. You can perform this test on a daily basis or as needed to verify the persistence of failures detected earlier.

This command verifies the functional operation of the switch by sending frames from a port's transmitter, and looping the frames back through an external fiber cable into the port's receiver. The test exercises all switch components from the main board, to the fibre cable, to the media (of the devices and the switch), and back to the main board.

The cables and media connected should be of the same type: a short-wavelength media (switch) port should be connected to another short-wavelength media (device) port using a short-wavelength cable; a long-wavelength port should be connected to a long-wavelength port, and a copper port should be connected to a copper port.

Only one frame is transmitted and received at any given time. The port LEDs flicker green while the test is running.

portTest

This command supports E_Ports, F_Ports (must support ELS Echo), L_Port, and N->N loopback ports. In addition, on switches running Fabric OS v6.4.0 and later, you can now use **portTest** on port configurations that previously caused nonspecific test results or were skipped by **portTest**. The following ports are now support the **portTest** diagnostics.

- E_Ports
- Static D_Ports
- N->N Loopback Ports
- F_Ports connected to Brocade-branded HBA
- Ports with index numbers greater than 255.
- Ports with swapped areas.
- Ports in shared area regions.
- Ports in logical switches.
- Ports in Base Switches.
- Long Distance ports.

This command is currently not supported on the following ports or switch configurations:

- Interchassis links (ICL) ports
- F_Ports connected to non-Brocade-branded HBA
- F_Ports connected to an Access Gateway
- F_Ports on VF-enabled switch
- EX_Ports
- E_Ports connected to EX_Ports
- Dynamic D_Ports
- On-demand D_Ports
- The **portTest** diagnostics is not supported in Access Gateway mode.

This command performs the following operations:

1. Initiates tests on certain ports (**portTest** command).
2. Stops active tests on certain ports (**stopPortTest** command).
3. Takes a snapshot of the test result (**portTestShow** command).

Use the **stopPortTest** command to stop the test. Refer to the **stopPortTest** help page for more information.

Use the **portTestShow** command to view the current status of **portTest**. Refer to the **portTestShow** help page for more information.

If there is a port type change during **portTest** execution, the test continues on a given port as long as it can be supported and it is asked to do so. If a request was made to test all ports on a given switch, **portTest** starts a new test appropriate for the new port type.

Notes The test is skipped if the F_Ports are connected to non-Brocade-branded devices. The test is also skipped on F_Ports if Virtual Fabric is enabled on any platform.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

- ports *itemlist*** Specifies the list of user ports to test. By default, all user ports in the switch are tested. Refer to the **itemList** help page for more information.
- iteration *count*** Specifies the number of times (or number of frames per port) to execute this test. Specify 0 to run the test in timeout mode, or specify -1 to run the test indefinitely. The default value is 20.
- userdelay *time*** Specifies the delay between frames sent by **portTest**, in milliseconds. The default value is 10 milliseconds.
- timeout *time*** Specifies the number of seconds to run the test. The timeout option goes into effect only if iteration is set to zero. The default value of timeout is 0.
- pattern *pattern*** Specifies the pattern of the test packet payload. The pattern is selected from the set of the first twenty predefined pattern types. Use the **dataTypeShow** command to view the patterns supported with **portTest**. For each pattern, the **dataTypeShow** command displays the name, the pattern type number, and an example. Specify the pattern by its type number. If **pattern** is not specified, it defaults to RANDOM (type=11).
- patsize *size*** Specifies the size of the pattern. The default pattern size is 1024 bytes. The range is 4 to 2048 bytes.
- seed *seed*** Specifies the seed value to be used with the pattern. The default seed value is Oxaa.
- listtype *porttype*** Specifies the type of ports on which to run **portTest**. Valid values for **porttype** include the following:
 - 1** All ports (default).
 - 2** All L_Ports.
 - 3** All F_Ports.
 - 4** All E_Ports.

portTest

-5

All N->N loopback ports.

Examples To run a functional test on an active switch:

```
switch:admin> porttest -ports 1/1-1/3
```

SEE ALSO [portLoopbackTest](#), [portTestShow](#), [spinFab](#), [stopPortTest](#)

portTestShow

Displays information from **portTest**.

Synopsis `porttestshow [-ports itemlist]`
`porttestshow -i | -index [index1[-index2]]`
`porttestshow -x [hex1[-hex2]]`

Description Use this command to display a snapshot of information from **portTest**. The command output displays statistical data about past test runs and values for parameters that were set when the test was run. Refer to the **portTest** help page for more information on the displayed parameters. The output includes the following information:

| | |
|-------------------------|--|
| Port number | Displays test status for the port. Values are PASS or FAIL. |
| PortType | Type of port tested. |
| PortState | Current State of portTest . Values are NO TEST, TESTING, or TEST DONE. |
| PortTypeToTest | Groups of ports tested. Values are ALL_PORTS, ALL_E_PORTS, ALL_L_PORTS, ALL_F_PORTS, ALL_LB_PORTS, or SINGLE_PORT. |
| Pattern | The pattern of the test packets payload. |
| Seed | The seed value used with the selected pattern. The default seed value is Oxaa. |
| UserDelay | The delay between frames sent by portTest in milliseconds. The default value is 10 milliseconds (ms). |
| TotalIteration | Total number of test iterations. |
| CurrentIteration | Current test iteration. |
| TotalFail | Total number of failed tests. |
| ConsecutiveFail | Number of consecutive failed tests. |
| StartTime | portTest start time. |
| StopTime | portTest end time. |

Timeout
Number of timed out tests.

ErrorCode
Error code, if any.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

- ports *itemlist***
Displays test results for the specified ports. This operand is optional; if omitted, data for all ports are displayed. Refer to the **itemList** help page for further details.
- index | -i *indexlist***
Displays test results for a single port or for a range of ports identified by port index numbers. You may specify index ranges separated by "-", for example, -i 33-47. This operand is optional, if omitted data for all ports are displayed.
- x [*indexhexlist*]**
Displays test results for a port or a range of ports identified by port index numbers in hexadecimal format. You may specify port ranges separated by "-", for example, -x 0x21-0x26. This operand is optional, if omitted data for all ports are displayed.

Examples To display information from **portTest**:

```
switch:admin> porttestshow 1
Port 1 : PASS
PortType: OTHER           PortState: NO TEST
PortInternalState: INIT   PortTypeToTest: NO_TEST
Pattern: 0x0      Seed: 0x0      UserDelay: 0

TotalIteration: 0           CurrentIteration: 0
TotalFail: 0               ConsecutiveFail: 0
StartTime: NONE
StopTime: NONE
Timeout: 0                 ErrorCode: 0
```

To display information from **portTest** on a port specified by the index number:

```
switch:admin> porttestshow -index 32
Port 32 : PASS
PortType: E PORT           PortState: TEST DONE
PortInternalState: INIT   PortTypeToTest: NO_TEST
Pattern: 0xb      Seed: 0xaa      UserDelay: 10

TotalIteration: 20           CurrentIteration: 20
TotalFail: 0               ConsecutiveFail: 0
StartTime: Wed Aug 13 06:39:59 2014
StopTime: Wed Aug 13 06:40:19 2014
Timeout: 0                 ErrorCode: 0
```

To display information from **portTest** on a port specified by the index number in hexadecimal format:

```
switch:admin> porttestshow -x 0x21
Port 33 : PASS
PortType: E PORT                      PortState: TEST DONE
PortInternalState: INIT                PortTypeToTest: NO_TEST
Pattern: 0xb                          Seed: 0xaa                      UserDelay: 10

TotalIteration: 20                    CurrentIteration: 20
TotalFail: 0                          ConsecutiveFail: 0
StartTime: Wed Aug 13 06:39:59 2014
StopTime: Wed Aug 13 06:40:19 2014
Timeout: 0                            ErrorCode: 0
```

SEE ALSO [portLoopbackTest](#), [portTest](#), [spinFab](#), [stopPortTest](#)

portTrunkArea

Assigns or removes a trunk area (TA) from a port or port trunk group; displays masterless F_Port trunking configuration.

Synopsis `porttrunkarea --enable [slot/port1[-port2]] -index port_index`

`porttrunkarea --disable [slot/port1[-port2]]`

`porttrunkarea --disable all`

`porttrunkarea --show disabled | enabled | trunk | all`

`porttrunkarea --show slot/port1[-port2]`

Description

Use this command to assign a static trunk area (TA) on a port or port trunk group, to remove a TA from a port or group of ports in a trunk, and to display masterless F_Port trunking information. The TA is identified by the port index number displayed in the output of the **switchShow** command.

Masterless F_Port trunking interoperates between the Access Gateway (AG) and Condor-based platforms. It is designed to (1) prevent reassignments of virtual addresses when F_Ports come back online after going offline and (2) to increase N_Port bandwidth.

Assigning a TA to a port or trunk group enables F_Port masterless trunking on that port or trunk group. When a TA is assigned to a port or trunk group, the ports immediately acquires the TA as the area of their process IDs (PID). Likewise, when a TA is removed from a port or trunk group, the ports reverts to the default area as their PID.

Use the **--show** option to obtain configuration details including the following information.

| | |
|---------------|--|
| Slot | On enterprise-class platforms, displays the slot number. |
| Port | Displays the port number. |
| Type | Displays online masterless trunked F_Port or EX_Port if applicable. Otherwise displays --. |
| State | Displays Trunk Master, Slave, or --. |
| Master | Displays the master port of the trunk group. |
| TA | On standalone switches, displays the user assigned TA number. |
| DA | On standalone switches, displays the default port area. The default area can be a port swapped area. |
| TI | On enterprise-class platforms, displays the user-assigned TA port index. |

| | |
|--|---|
| DI | On enterprise-class platforms, displays the default port index. The default port index can be a port swapped area. |
| The --show trunk option displays the following information: | |
| Trunk Index | Displays the trunk index. |
| ptA->ptB | ptA indicates the local user port; ptB indicates the remote user port. |
| sp | Port speed in Gbps. |
| Bandwidth | The bandwidth (Rx, Tx, and the combined total for Tx+Rx) of the trunk group. Values are displayed as either bits per second (Bps), kilobits per second (Kbps), megabits per second (Mbps), or gigabits per second (Gbps), rounded down to the next integer. |
| Throughput | Displays the throughput (Rx, Tx, and the combined total for Tx+Rx) of the trunk group. Results are displayed for the previous second. Values are displayed as either bits per second (Bps), kilobits per second (kbps), megabits per second (Mbps), or gigabits per second (Gbps), rounded down to the next integer. |
| % | Displays the percentage of link utilization (Rx, Tx, and the combined total for Tx+Rx) . Even when the link utilization is 100% , the throughput value will be lesser than the bandwidth value, due to the 8b/10b encoding and the control words transmitted. For example, the throughput for an 8 Gbps link at 100% utilization would be approximately 6.8 Gbps. |
| deskew | The time difference for traffic to travel over each F_Port trunk as compared to the F_Port trunk with the shortest travel time in the group. The value is expressed in nanoseconds divided by 10. The firmware automatically sets the minimum deskew value of the shortest F_Port trunk travel time to 15. |
| Master | Identifies the master port of the trunk group. |

Execution of this command is subject to the following restrictions:

- Only F_Port trunk ports are allowed to be part of a TA. E/F/L/EX_Port will be persistently disabled. Private L_Ports remain online but will not run traffic.
- Only one trunk master per TA is permitted. The second trunk master is persistently disabled.
- The entire TA trunk group shares the same port WWN.
- The port must be disabled before a TA can be assigned to a port or removed from a trunk group.
- There is one port whose Default Area is the same as its Trunk Area. You cannot remove that port from the trunk group unless The TA is removed from all ports in the trunk group.
- You must enable trunking on all ports to be included in a TA before you can create a TA. Use **portCfgTrunkPort** or **switchCfgTrunk** to enable Trunking on a port or on all ports of a switch.
- N_Port ID Virtualization (NPIV) support is provided for up to 255 devices per TA. Note that this decreases the number of devices available per port, because all participating ports share the same area.
- F_Port trunking is only supported in CORE PID formats.
- Certain port configuration features are not supported within a TA and the command fails if one of these features is enabled on a port within the trunk group. These features include FastWrite, Port Swapping, Port Mirroring, Long Distance, Interchassis links (ICL), and FICON.
- Ports included in a TA share the same port index. The original port index may be removed in the process. This means that D, I zones referring to these indices are no longer part of the switch. For details and workarounds, refer to the Fabric OS Administrator's Guide.
- Device Connection Control (DCC) Policy must be removed from ports prior to creating a TA. You can re-enable DCC policy after creating the TA.
- You cannot assign a TA while AG mode is enabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Masterless F_Port trunking requires an ISL Trunking license.

Operands This command supports the following operands:

slot

On enterprise-class platforms, specifies the slot number, followed by a slash (/).

port1[-port2]

Specifies a single port or a port range, relative to its slot on bladed systems. For example, 9/8-15 on an enterprise-class platform indicates slot 9, ports 8 to 15. Port ranges should fall in the octet (8 port) trunk range starting from port 0 on a switch or blade.

--enable

Creates a TA assigned to the specified ports. Use this option with one of the following operands:

-area *area_number*

On single processor switches, specifies the port area number for the static TA to be created. The TA must fall within the 8 port trunk group starting from port 0 on a switch or blade. The TA must be a default area of an existing port in a given port group. However, you may add ports to the trunk group even when a trunk group has already been assigned by using the same area of the octet trunk group. Use **switchShow** for a list of valid port area or index numbers.

-index *port_index*

On enterprise-class platforms, specifies the port index for the static TA to be created. The port index must fall within the 8 port trunk group starting from port 0 on a switch or blade. The TA must be a default index of an existing port in a given port group. However, you may add ports to the trunk group even when a trunk group has already been assigned by using the same index of the octet trunk group. Use **switchShow** for a list of valid port indexes.

--disable

Removes specified ports from a TA. If a port with the same default area as the TA assigned for the trunk group is removed, all ports in the trunk group must be explicitly specified for removal.

all

Optionally removes all TA assigned ports on the switch. This option disables masterless F_Port trunking on all ports. All TA assigned ports must be disabled for this option to succeed.

--show

Displays masterless F_Port trunking information. When using this option, specify one of the following operands:

[*slot*]/*port1*[-*port2*]

Displays configuration for a specified port or port range.

trunk

Displays configuration details for the port trunk group, including user port, neighboring user port, and master port properties.

enabled

Displays configuration details for all ports included in a user assigned TA (all ports on which masterless F_Port trunking is enabled).

disabled

Displays configuration details for all ports not included in a user assigned TA (all ports on which masterless F_Port trunking is not enabled).

all

Displays configuration details for all ports on a switch.

Examples

To enable masterless F_Port trunking on a standalone switch:

1. Disable ports 10-11 by executing **portdisable *port*** for each port to be included in the TA.
2. Enable Trunk Area for ports 10-11 with area number 37:

```
switch:admin> porttrunkarea --enable 10-11 -index 11
2009/05/15-12:43:10, [SWCH-1012], 60, FID 128, INFO,
sw0, Trunk Area (11) has been enabled for one or more ports
Trunk area 11 enabled for ports 10 and 11.
```

3. Re-enable ports 10-11 by executing **portenable *port*** for each port in the TA.
4. Show switch/port information:

```
switch:admin> switchshow
[...]
Index Port Address Media Speed State Proto
=====
[...]
11 10 030b00 id N4 No_Light FC
11 11 030b00 id N4 No_Light FC
[...]
```

5. Display TA-enabled port configuration:

```
switch:admin> porttrunkarea --show enabled
```

```
Port Type State Master TA DA
-----
10 -- -- -- 11 10
11 -- -- -- 11 11
```

To disable masterless F_Port trunking on ports 10-11:

```
switch:admin> porttrunkarea --disable 10-11
ERROR: port 11 has to be disabled
```

Disable each port prior to removing ports from the TA. Then reissue the command:

```
switch:admin> porttrunkarea --disable 10-11
Trunk area 11 disabled for ports 10 and 11.
```

To display trunk details for a user assigned TA 25 that includes ports 24-25:

```
switch:admin> porttrunkarea --show trunk
Trunk Index 25: 25->0 sp: 8.000G \
bw: 16.000G deskew 15 MASTER
Tx: Bandwidth 16.00Gbps, Throughput 1.63Gbps (11.84%)
Rx: Bandwidth 16.00Gbps, Throughput 1.62Gbps (11.76%)

Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.24Gbps (11.80%)
24->1 sp: 8.000G bw: 8.000G deskew 15
Tx: Bandwidth 16.00Gbps, Throughput 1.63Gbps (11.84%)
Rx: Bandwidth 16.00Gbps, Throughput 1.62Gbps (11.76%)
Tx+Rx: Bandwidth 32.00Gbps, Throughput 3.24Gbps (11.80%)
```


To configure a TA on an enterprise-class platform including ports 13 and 14 on slot 10 with port index of 125:

1. Disable the ports to be included in the TA.
2. Enable TA for ports 13 and 14 on slot 10 with port index of 125:

```
switch:admin> porttrunkarea --enable 10/13-14 -index 125
Trunk index 125 enabled for ports 10/13 and 10/14.
```

3. Show the TA port configuration (ports still disabled):

```
switch:admin> porttrunkarea --show enabled
Slot Port Type State Master TI DI
-----
10 13 -- -- -- 125 125
10 14 -- -- -- 125 126
-----
```

4. Enable ports 13 and 14:

```
switch:admin> portenable 10/13
switch:admin> portenable 10/14
```

5. Show the TA port configuration after enabling the ports:

```
switch:admin> porttrunkarea --show enabled
Slot Port Type State Master TI DI
-----
10 13 F-port Master 10/13 125 125
10 14 F-port Slave 10/13 125 126
```

SEE ALSO [portCfgTrunkPort](#), [portCfgShow](#), [portShow](#), [switchCfgTrunk](#), [switchShow](#)

portZoneShow

Displays the enforced zone type of the F_Ports and FL_Ports of a switch.

Synopsis `portzoneshow`

Description Use this command to display the enforced zone type of the F_Ports and FL_Ports of a switch.

Output shows virtual port number (decimal), physical port number (decimal), online status, and if online, port type. If the current zone configuration has been disabled by **cfgDisable**, the fabric is in non-zoning mode, in which all devices see each other. When default zoning is enabled with "No Access" mode, "No Effective configuration: (No Access)" is displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the zone membership information of ports:

```
switch:user> portzoneshow
[OUTPUT TRUNCATED]
PORT: 160 (160)  Offline
PORT: 161 (161)  Offline
PORT: 162 (162)  Offline
PORT: 163 (163)  Offline
PORT: 164 (164)  Offline
PORT: 165 (165)  Offline
PORT: 166 (166)  Offline
PORT: 167 (167)  Offline
PORT: 168 (168)  FL-Port  Enforcement: HARD WWN \
                    defaultHard: 0  IFID: 0x4332000a
PORT: 169 (169)  Offline
PORT: 170 (170)  Offline
PORT: 171 (171)  Offline
PORT: 172 (172)  Offline
PORT: 173 (173)  Offline
PORT: 174 (174)  Offline
PORT: 175 (175)  Offline
PORT: 176 (176)  F-Port  Enforcement: HARD WWN  \
                    defaultHard: 0  IFID: 0x4342002a
PORT: 177 (177)  Offline
PORT: 178 (178)  Offline
PORT: 179 (179)  Offline
PORT: 180 (180)  Offline
(output truncated)
```

SEE ALSO [cfgShow](#), [switchShow](#)

powerOffListSet

Sets the order in which slots are powered off.

Synopsis `powerofflistset`

Description Use this command to modify the order in which slots are powered off. This command displays the current order, and then prompts you interactively to confirm or modify the power-off position for each slot.

Whenever a power supply goes out of service or a field-replaceable unit (FRU) is inserted, the system's available power is compared to the system's required power to determine if there is enough power to operate. If less than the required power is available, the power-off list is processed, until there is sufficient power for the system to operate.

If the system's power supply drops abruptly to insufficient levels, the power-off list cannot be processed. The sudden lack of power causes the CP board processors to cease executing the firmware.

For example, if only one power supply is available to power a fully loaded system and the power supply is removed from the chassis, all system operations terminate immediately, and the power-off list cannot be processed. However, if the system is running on two power supplies (this is not recommended) and one goes into a predicted fail state (in which the power supply is still supplying power), the power-off list is processed as described.

The power-off list does not affect the order, in which slots are powered on. On power-on or when an additional power supply is added, slots are processed sequentially, starting at slot 1.

Notes CP blade slots and Core blade slots are not included in the power-off list.

Command output may vary depending on the hardware platform.

Some FRUs may use significant power, but cannot be powered off by the software. For example a missing blower FRU may change the power computation enough to affect how many slot blades can be powered up.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To modify the power-off list order:

```
switch:admin> powerofflistset

Slot      Current POL
-----
12         1st
11         2nd
10         3rd
9          4th
6          5th
5          6th
4          7th
3          8th

1st slot to be power off: (3..12) [12] 3
```

powerOffListSet

```
2nd slot to be power off: (4..12) [11] 4
3rd slot to be power off: (5..12) [10] 5
4th slot to be power off: (6..12) [9] 6
5th slot to be power off: (9..12) [9] 9
6th slot to be power off: (10..12) [10] 10
7th slot to be power off: (11..12) [11] 11
8th slot to be power off: (12..12) [12] 12
```

| Old POL | New POL | Power Off Order |
|---------|---------|-----------------|
| 12 | 3 | 1st |
| 11 | 4 | 2nd |
| 10 | 5 | 3rd |
| 9 | 6 | 4th |
| 6 | 9 | 5th |
| 5 | 10 | 6th |
| 4 | 11 | 7th |
| 3 | 12 | 8th |

```
Proceed to change the POL order? (yes, y, no, n): [no] y
```

SEE ALSO [chassisShow](#), [powerOffListShow](#), [psShow](#), [slotPowerOff](#), [slotPowerOn](#), [slotShow](#)

powerOffListShow

Displays the order in which slots are powered off.

Synopsis `powerofflistshow`

Description Use this command to display the order in which the physical slots are powered off.

Whenever a power supply goes out of service or a field-replaceable unit (FRU) is inserted, the system's available power is compared to the system's required power to determine if there is enough power to operate. If less than the required power is available, the power-off list is processed, until there is sufficient power for the system to operate.

The following rules apply when the power-off list is processed:

- If a power supply starts predicting failure, so that there will not be enough power for all blades, the powered-up port blades are powered down in the order in which they appear on the power-off list.
- If you replace a failed power supply, or you remedy an existing power-supply shortage by inserting a new power supply, the previously powered-down blades are powered up in the reverse order of the power-on list. Note that this is different from the order in which the blades in the system usually come up. The original powering up of blades does not consult the power-off-list; it simply proceeds from the lower-numbered slots to the higher-numbered slots.
- If you add a new blade to the chassis, and there is not enough power available to operate the additional hardware (because you may only have one power supply) the newly inserted blade will be denied power and the existing powered-on blades stay powered on.
- If the system's power supply drops abruptly to insufficient levels, the power-off list cannot be processed. The sudden lack of power causes the CP board processors to cease executing the firmware. For example, if only two power supplies are available to power a fully loaded chassis that requires at least two power supplies, and one power supply is removed from the chassis, all system operations terminate immediately, and the power-off list cannot be processed. However, if the system is running on two power supplies (this is not recommended in this case) and one goes into a predicted fail state (in which the power supply is still supplying power), the power-off list is processed as described.

Notes Command output may vary depending on the hardware platform.

Control processor (CP) blades are not included in the power-off list.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the slot power off list order:

```
switch:admin> powerofflistshow

Slot 12 will be powered off 1st
Slot 11 will be powered off 2nd
Slot 10 will be powered off 3rd
```

powerOffListShow

```
Slot 9 will be powered off 4th  
Slot 4 will be powered off 5th  
Slot 3 will be powered off 6th  
Slot 6 will be powered off 7th  
Slot 5 will be powered off 8th
```

SEE ALSO [chassisShow](#), [powerOffListSet](#), [psShow](#), [slotPowerOff](#), [slotPowerOn](#), [slotShow](#)

psShow

Displays power supply status.

Synopsis `psshow`

Description Use this command to display the current status of the switch power supplies.

The status of each supply is displayed as:

OK

Power supply functioning correctly.

absent

Power supply not present.

unknown

Unknown power supply unit installed.

predicting failure

Power supply is present but predicting failure.

faulty

Power supply present but faulty (no power cable, power switch turned off, fuse blown, or other internal error).

For certain switch models, the OEM serial ID data displays after each power supply status line.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To view the status of the power supplies:

```
switch:admin> psshow

Power Supply #1 is OK
  DELTA DPS-1001AB-1E 23000000601 S1   IXD0111000088
Power Supply #2 is faulty
  DELTA DPS-1001AB-1E 23000000601 S1   IXD0111000162
Power Supply #3 is OK
  DELTA DPS-1001AB-1E 23000000601 S1   IXD0111000120
Power Supply #4 is absent
```

SEE ALSO [chassisShow](#), [fanShow](#)

rasAdmin

Configures RASLog message generation.

Synopsis `rasadmin --enable [-log MSG-ID] [-syslog MSG-ID]
[-module MODULE-ID]`

`rasadmin --disable [-log MSG-ID] [-syslog MSG-ID]
[-module MODULE-ID]`

`rasadmin --set -log MSG-ID -severity value`

`rasadmin --show [-log MSG-ID] [-severity MSG-ID]
[-module MODULE-ID] [-disabled] [-syslog] [-all]`

`rasadmin --help`

Description Use this command to enable or disable RASLog message logging for selected messages or groups of messages (modules), to change the default severity level for a specified message, and to display configured RASLog settings. The **-log**, **-module**, and **-severity** configuration and display options apply to external messages (Message ID 1001–4999) and are available to any user with admin privileges.

An INFO RASLog message is generated for every message that is enabled or disabled. In addition, the list of disabled RASLog messages are collected as part of the **supportSave** command

The changes made by this command are persistent across reboots, high availability failover, and firmware downloads.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command does not disable messages of type FFDC or AUDIT. When you disable logging for a message that has both AUDIT and LOG attributes, the message will not be logged as a RASLog message, but it will continue to be written to the Audit log.

Operands This command has the following operands:

--disable

Disables logging for a specified RASLog message or for all messages grouped in a RASLog module.

--enable

Enables logging for a specified RASLog message or for all messages grouped in a RASLog module. Message logging is enabled by default. This command re-enables logging of messages that were previously disabled.

One of the following options must be specified when you disable or enable message logging:

-log *MSG-ID*

Enables or disables logging for the specified message.

- log *MODULE-ID***
Enables or disables logging for all message included in the specified module.
- syslog *MSG-ID***
Enables or disables logging for an internal message.
- set -log *MSG-ID***
Changes the default severity level of the specified message.
- severity *value***
Specifies a new severity level for the message. Valid values include **INFO**, **WARNING**, **ERROR**, **CRITICAL**, and **DEFAULT**.
- show -disabled**
Displays all messages that have been disabled.
- show -log *MSG-ID***
Displays the logging status of the specified message.
- show -module *MODULE-ID***
Displays the logging status of all messages included in the specified module.
- show -severity *MSG-ID***
Displays the severity of the specified messages.
- show -all**
Displays all external RASlog messages, their status (enabled or disabled), their configured severity and their default severity.
- show -syslog**
Displays the list of internal RASLog messages that are configured to display on the system console and the syslog server.

Examples

To disable logging of a single message:

```
switch:admin> rasadmin --disable -log NSM-1009
2012/07/20-13:30:41, [LOG-1005], 378, SLOT 4 | CHASSIS, INFO,
PLUTO_25, Log message NSM-1009 has been disabled.
```

To re-enable logging of a single message that was previously disabled:

```
switch:admin> rasadmin --enable -log NSM-1009
2012/07/20-13:30:41, [LOG-1005], 378, SLOT 4 | CHASSIS, INFO,
PLUTO_25, Log message NSM-1009 has been enabled.
```

To disable logging of all messages that belong to the NSM module:

```
switch:admin> rasadmin --disable -module NSM
2012/07/20-13:28:37, [LOG-1007], 375, SLOT 4 | CHASSIS, INFO,
PLUTO_25, Log Module NSM has been disabled.
```

To re-enable logging of all messages that belong to the NSM module:

```
switch:admin> rasadmin --enable -module NSM
2012/07/20-13:28:37, [LOG-1007], 375, SLOT 4 | CHASSIS, INFO, PLUTO_25,
Log Module NSM has been enabled,
```

To change the severity level of a RASlog message:

```
switch:admin> rasadmin --set -log SEC-1203 -severity WARNING
Message      Severity
SEC-1203:    WARNING
```

To display a list of all messages that have been disabled:

```
switch:admin> rasadmin --show -disabled
Message      Status   Default Severity   Current Severity
IPAD-1002    DISABLED INFO
IPAD-1003    DISABLED INFO
```

To display the status and configuration of messages that belong to the specified module:

```
switch:admin> rasadmin --show -module RM
Message      Status   Default Severity   Current Severity
RM-1001      ENABLED  INFO
RM-1023      ENABLED  INFO
RM-1024      ENABLED  INFO
RM-1020      ENABLED  INFO
RM-1021      ENABLED  INFO
RM-1022      ENABLED  INFO
```

To display the status and configuration of a specified message.

```
switch:admin> rasadmin --show -log IPAD-1002
Message      Status   Default Severity   Current Severity
IPAD-1002    DISABLED INFO
```

To display the status and configuration for all external messages:

```
switch:admin> rasadmin --show -all
Message      Status   Default Severity   Current Severity
FCIP-1000    ENABLED  CRITICAL          CRITICAL
FCIP-1001    ENABLED  INFO              ERROR
FCIP-1002    ENABLED  INFO              INFO
```

To enable an internal RASlog messages to be sent to syslog (this is done per instruction from support):

```
switch:admin> rasadmin --enable -syslog RAS-5001
2012/07/20-13:49:35, [LOG-1009], 385, SLOT 4 | CHASSIS, INFO, PLUTO_25,
Internal Log message RAS-5001 has been enabled for syslog logging.
```

To disable an internal RASlog messages to be sent to syslog (this is done per instruction from support):

```
switch:admin> rasadmin --disable -syslog RAS-5001
2012/07/20-13:49:35, [LOG-1010], 385, SLOT 4 | CHASSIS, INFO, PLUTO_25,
Internal Log message RAS-5001 has been disabled for syslog logging.
```

SEE ALSO [rasMan](#)

rasMan

Displays RASlog message text and documentation.

Synopsis `rasman message_id`

`rasman --help`

Description Use this command to display documentation for a specified RASlog message. The message is specified by its message ID. The command output is identical to the documentation provided in the *Fabric OS Message Reference*. For each message, the command displays the following information:

MESSAGE

Displays the message text.

MESSAGE TYPE

Displays the message type. The message type can be one or more of the following: LOG, FFDC, or AUDIT.

CLASS

A class value is displayed for audit messages only. Supported audit classes include SECURITY, FIRMWARE, FABRIC, LS (logical switch), CLI, MAPS, and RAS.

SEVERITY

Indicates the message severity as one of the following: INFO, ERROR, WARNING, CRITICAL.

PROBABLE CAUSE

Describes what may be causing the message.

RECOMMENDED ACTION

Describes the recommended action.

Notes This command is available to all supported Brocade roles, including root.

This command displays only external messages in the numeric range of 1000-4999.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

message_id

Specifies the message ID in the following format: <module ID>-<message number>. The message ID is case sensitive and should be entered exactly as shown in the RASlog message displayed on the console. For example, ZONE-3018 is a valid message ID, but zone-3018 is not valid.

`--help`

Displays the command usage.

Examples To display documentation for a ZONE audit message:

```
switch:admin> rasman ZONE-3018
Log Messages                                     ZONE-3018 (7m)

MESSAGE
  ZONE-3018 - Event: <Event Name> Status:  success,
  Info: <AD object type> <AD object name>
  has been deactivated.

MESSAGE TYPE
  AUDIT

CLASS
  FABRIC

SEVERITY
  INFO

PROBABLE CAUSE
  Indicates that the specified Admin Domain (AD) object has
  been deactivated.

RECOMMENDED ACTION
  Verify that the event was planned. If the event was planned,
  no action is required. If the event was not planned, take
  appropriate action as defined by your enterprise security
  policy.
```

SEE ALSO [rasAdmin](#)

reboot

Reboots the control processor (CP).

Synopsis `reboot [-f]`

Description Use this command to perform a "cold reboot" (power off/restart) of the control processor. This operation may be disruptive, and the command prompts for confirmation before executing. When you reboot a switch connected to a fabric, all traffic to and from that switch stops. All Fibre Channel ports on that switch including E_Ports become inactive until the switch comes back online.

The behavior of this command depends on the platform:

- When issued on a standalone (single-processor) switch, this command performs a cold reboot of the switch.
- When issued on an enterprise-class platform (Brocade X6-4 Director, Brocade X6-8 Director, DCX 8510-4, or DCX 8510-8) with two CPs (active and standby), the following rules apply:
 - When the Standby CP reboots, it goes down and there is no failover because there is no traffic on that switch. When the Standby CP comes up again, it is temporarily no longer in sync with the Active CP.
 - When the Active CP reboots, it fails over to the Standby CP. The Standby CP becomes the new Active CP and traffic is disrupted.
 - When HA is in sync, and **reboot -f** is issued on the Active CP of a director, the Standby CP takes over as the active CP without traffic disruption. If HA is not in sync, and **reboot -f** is issued on the Active CP, the Standby CP takes over as the Active CP and traffic is disrupted.
 - When HA is disabled and **reboot** or **reboot -f** is issued on the Active CP, both the Active and Standby CPs reboot with the original mastership retained. The original Active CP remains the Active CP after the reboot, and the original Standby CP remains the Standby CP. After the reboot, HA is enabled.
 - When HA is disabled and **reboot** or **reboot -f** is issued on the Standby CP, the Standby CP reboots without prompting. It boots up with the default switch only, even if the Active CP has multiple logical switches configured. After the Standby CP boots up, HA is still disabled.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands `-f` Causes the CP to fail over to the Standby CP without affecting any of the Fibre Channel Ports. Note that the recommended way to force a failover without affecting any Fibre Channel ports is to issue the **haFailover** command.

Examples To reboot a standalone switch with a single CP:

```
switch:admin> reboot
Warning: This command would cause the switch to reboot
and result in traffic disruption.
Are you sure you want to reboot the switch [y/n]? y

Broadcast message from root (pts/0) Sun Feb 28 19:49:45 2010...
The system is going down for reboot NOW !!
```

reboot

To reboot a CP on a DCX when HA is enabled:

```
switch:admin> reboot
Warning: This command is being run on a control processor (CP)
based system and will cause the active CP to reboot.
Are you sure you want to reboot the active CP [y/n]? y

Broadcast message from root (pts/0) Sun Feb 28 19:49:45 2010...

The system is going down for reboot NOW !!
```

To reboot a CP on a DCX when **haFailover** is disabled:

```
switch:admin> reboot

This command is being run on a control processor (CP)
based system. Because HA is disabled, it will cause both
active CP and the standby CP to reboot. After reboot, the
HA will be enabled.

Do you want to continue [y/n] y

Broadcast message from root (pts/0) Sun Feb 28 19:49:45 2010...

The system is going down for reboot NOW !!
```

SEE ALSO [fastBoot](#), [haDisable](#), [haEnable](#), [haFailover](#)

relayConfig

Sets and displays the relay host IP address.

Synopsis `relayconfig --config -rla_ip relay_ip
-rla_dname domain_name`

`relayconfig --delete`

`relayconfig --show`

`relayconfig --help`

Description Use this command to configure or display the relay host that is used to send the Monitoring and Alerting Policy Suite (MAPS) e-mail notifications. You can configure the e-mail recipients using the `mapsConfig --emailcfg` command.

For an e-mail alert to function correctly, add the CPO and CP1 IP addresses and host names to DNS and also set the domain name and name server. The `ipAddrShow` and `dnsConfig` commands can be used to set and verify this information.

Notes This command requires a Fabric Vision license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--config Sets the relay host IP address and domain name.

`-rla_ip relay_ip` Specifies the IP address of the relay host.

`-rla_dname domain_name` Specifies the domain name of the relay host.

--delete Deletes the relay host configuration.

--show Displays the relay host configuration.

--help Displays the command usage.

Examples To set relay host IP address:

```
switch:admin> relayconfig --config -rla_ip 10.70.212.168 \  
-rla_dname "mail.brocade.com"
```

relayConfig

To display the relay host configuration:

```
switch:admin> relayconfig --show
Relay Host:          10.70.212.168
Relay Domain Name: mail.brocade.com
```

To delete the relay host configuration:

```
switch:admin> relayconfig --delete
```

SEE ALSO [mapsConfig](#)

roleConfig

Manages user-defined roles.

Synopsis

```
roleconfig --add role_name [-desc description]
  [-class rbac_class_list] [-perm permission]

roleconfig --change role_name [-class rbac_class_list
  -perm permission [-desc description]

roleconfig --delete role_name [-force]

roleconfig --copy new_role -role source_role

roleconfig --show role_name | -all [default]

roleconfig --help
```

Description Use this command to create or modify user-defined roles, to define permissions for these roles based on role-based access control (RBAC) permissions and meta-object format (MOF) classes, and to display the configured roles. Two types of access control restriction exist in Fabric OS:

- Restriction by MOF class: A MOF class groups similar Fabric OS commands into feature sets that share the same access permissions. By assigning one or more MOF classes to a role, the account with the specified role can access all the commands included in these classes. For example, the predefined role ZoneAdmin can access the commands under the MOF class Zoning, but not those under the UserManagement class. With the **roleConfig** command you could define a special admin role called myzonesec and assign access to this role for both the zoning and the userManagement class.
- Restriction by RBAC access level: You can further restrict access by setting RBAC one of the following access levels for the role. The RBAC permissions are set per class.
 - O = observe
 - OM = observe-modify
 - N = none/not available

Use the **--show** option to display information about user-defined roles and default roles. Use the **classConfig** command to display information about MOF classes and associated commands. Note that you cannot modify the predefined Fabric OS roles.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--add

Creates a role with the specified name and optional attributes. The new role is created with two default RBAC classes, "localuserenvironment" and "nocheck" and has the default permissions observe and modify ("OM"). A configuration download will always reset the permissions of these two default classes to "OM."

--change

Modifies an existing user-defined role.

- role_name*
Specifies the name for the role to be created or modified. The name must be unique; it is case-insensitive and can contain only alpha characters. The role name must be at least 4 characters long and cannot exceed 16 characters. The maximum number of user-defined roles allowed on a chassis is 150. This operand is required.
- The following operands are optional with **--add** and **--change**:
- desc** *description*
Specifies a description for the role of up to 63 characters. Colons (:) are not permitted.
- class** *class_list*
Specifies one or more MOF classes to which the role should have access. Classes must be separated by commas. The requested class permissions cannot be higher than those of the Fabric OS Admin role.
- perm** *permissions*
Specifies the RBAC permissions for the role. The RBAC permissions restrict what the user can do with the commands included in the classes to which the role has access. Valid RBAC permissions include the following:
- OM** Observe and modify.
 - O** Observe only.
 - N** No access. This parameter is not valid with the **--add** option.
- delete** *role_name* [**-force**]
Deletes the specified user-defined role. This command prompts for confirmation unless you use the **-force** option. The role must exist in the database and the role cannot currently be assigned to a user account. You cannot delete any of the predefined Fabric OS roles.
- copy**
Clones an existing user-defined role by copying an existing role to a new role name. The new role inherits all the classes and permissions of the source role. You can further modify the newly created role. The following operands are required:
- new_role*
The name for the new role must be unique; it is case-insensitive and can contain only alpha characters. The role name must be at least 4 characters long and cannot exceed 16 characters.
- role** *source_role*
Specified the name of the exiting role to be copied.
- show**
Displays information about the specified roles. For each role, the command displays the role name, description, assigned classes and RBAC permissions for each class. The following displays options are exclusive:
- role_name*
Displays information about the specified user-defined role.

-all [default]

Displays a listing of all user-defined roles. When used together with the optional **-default** option, both the user-defined roles and the predefined Fabric OS roles are displayed.

--help

Displays the command usage.

Examples

To create a role with Zoning and SecurityAdmin permissions and OM access:

```
switch:admin> roleconfig --add myzonesec \
  -d "Zone and Security Admin" -c Security,zoning -p OM
```

To change the access permissions of the previously created role and to display the results:

```
switch:admin> roleconfig --change myzonesec -c pki -p N
switch:admin> roleconfig --change myzonesec -c security -p O
switch:admin> roleconfig --change myzonesec -c zoning -p OM

switch:admin> roleconfig --show myzonesec
Role Name      : myzonesec
Description    : Zone and Security Admin

RBAC Class                Permission
-----
Security                  O
Zoning                    OM

Role permissions changed successfully.
```

To copy the role myzonesec:

```
switch:admin> roleconfig --copy superrole -r myzonesec
Role 'superrole' is copied successfully.
```

To delete a role:

```
switch:admin> roleconfig --delete -r myzonesec
You are going to delete a user defined role.
Are you sure? (yes, y , no, n) [no] y
Role 'myzonesec' is deleted successfully.
```

SEE ALSO [classConfig](#)

rootAccess

Configures root access through CLI.

Synopsis `rootaccess --set [none | consoleonly | all]`

`rootaccess --show`

`rootaccess --help`

Description Use this command to display and set root access settings.

Use the **--set** command to configure the root user to login to the switch through console or other interfaces such as SSH, Telnet or Web Tools, or to block root login from the interfaces. By default, access to the root account is allowed through console only. Root account must be enabled for the changes to take effect. If root account is disabled, root login through console or any interface fails with a warning message. By default, root account is disabled. You can enable or disable root account using the **userConfig** command. When access is changed, root is logged out.

Use the **--show** command to display the current root access settings.

Notes The execution of this command is not subject to Virtual Fabric restrictions. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following optional operands:

| | |
|--------------------|--|
| --set | Modifies the root access settings. Valid options include the following: |
| none | Disables logging in to the switch as root through all interfaces. |
| consoleonly | Allows logging in to the switch as root through console only. This is the default setting. |
| all | Allows logging in to the switch as root through all interfaces. |
| --show | Displays the current root access settings. |
| --help | Displays the command usage. |

Examples To set root access level to 'all':

```
switch:admin> rootaccess --set all
```

To display root access settings:

```
switch:admin> rootaccess --show  
RootAccess: all
```

SEE ALSO [userConfig](#)

routeHelp

Displays a list of FSPF-related commands.

Synopsis `routehelp`

Description Use this command to display a list of fabric-shortest-path-first (FSPF)-related commands.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display a list of routing-related commands:

```
switch:admin> routehelp

aptPolicy          Get and set Advanced Performance
                   Tuning policy
bcastShow          Print broadcast tree information
dlsReset           Turn off Dynamic Load Sharing
dlsSet             Turn on Dynamic Load Sharing
dlsShow            Print state of Dynamic Load Sharing
fspfShow           Print FSPF global information
interfaceShow      Print FSPF interface information
iodReset           Turn off In-Order Delivery
iodSet             Turn on In-Order Delivery
iodShow            Print state of In-Order Delivery
linkCost           Set or print the FSPF cost of a link
LSDbShow           Print Link State Database entry
nbrStateShow       Print neighbor's summary information
nbrStatsClear      Reset FSPF neighbor's counters
topologyShow       Print paths to domain(s)
uRouteShow         Print port's unicast routing info
```

SEE ALSO [bcastShow](#), [interfaceShow](#), [uRouteShow](#)

rtLogTrace

Manages real-time trace logging.

Synopsis `rtlogtrace --enable`

`rtlogtrace --disable`

`rtlogtrace --show`

`rtlogtrace --help`

Description Use this command to enable or disable real-time trace logging and to indicate whether the feature is enabled or disabled.

The RTLog facility captures the most recent hardware events and low-level software interrupts on Brocade DCX platforms in real time. When enabled, the RTLog traces are collected as part of the **supportSave** utility. The RTLog generates no console output. The RTLog is enabled by default and persistent across reboots. The rtLogTrace configuration is lost after a power cycle.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------|--|
| --enable | Enables the RTLog on the active Control Processor (CP). The RTLog becomes effective on the standby CP after a failover or a reboot of the active CP. |
| --disable | Disables the RTLog. |
| --show | Displays the current status of the RTLog as enabled or disabled. |
| --help | Displays the command usage. |

Examples To enable the RTLog:

```
Switch:admin> rtlogtrace --enable
```

To disable the RTLog:

```
Switch:admin> rtlogtrace --disable
```

To display the RTLog status:

```
Switch:admin> rtlogtrace --show
```

SEE ALSO `None`

sddQuarantine

Clears or displays the quarantined ports.

Synopsis `sddquarantine --clear [[slot/]port | all] [-force]`

`sddquarantine --show`

`sddquarantine --help`

Description Use this command to clear or display ports that are quarantined by the MAPS action "SDDQ".

The slow drain device quarantine (SDDQ) feature is used to automatically detect the slow-drain devices and move the slow-drain flows to a low priority virtual channel (VC) from the existing VC (medium or high) thus freeing up the resources for the regular flows in the existing VC. Due to this automatic isolation from the regular flows, the effects of the slow-drain flows on the fabric are reduced. Once the traffic flowing through an F_Port is isolated, the F_Port is marked as Slow Drain Quarantined. You can use the **--clear** option to clear the Slow Drain Quarantined state on the port.

If the quarantined ports go offline or disabled, the ports remain in Slow Drain Quarantined state. Once the ports come online, the flows destined to the port are quarantined.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--clear

Clears the specified port or all ports from the Slow Drain Quarantined state. This option is not allowed if the latency or frame loss condition persists on the ports. The **-force** option overrides the default behavior. Valid options include the following:

slot

For bladed systems only, specifies the slot number for which to set or display the cost, followed by a slash (/).

port

Specifies the port number for which to set or display the cost, relative to its slot for bladed systems. Use **switchShow** to list of valid ports.

all

Clears all the ports currently quarantined in the local switch.

-force

Clears the quarantined ports even when the latency or frame loss condition persists. This operand is optional.

--show

Displays all the local quarantined ports, ports on which VC translation is not enforced due to zoning count restriction, and the fabric-wide quarantined device information.

--help

Displays the command usage.

Examples To clear quarantined state for a port:

```
switch:admin> sddquarantine --clear 33
Initiated clearing port from quarantined state
```

To clear quarantined state for all ports:

```
switch:admin> sddquarantine --clear all
The clear action was not initiated for the following port(s). Try with individual
ports
          3
Initiated quarantine action on other ports
```

To display the offline quarantined local ports and the online quarantined device information across the fabric:

```
switch:admin> sddquarantine --show
-----
Ports marked as Slow Drain Quarantined in the Local Switch: 2/1, 1/3
-----
Ports marked as Slow Drain Quarantined but not enforced:    1/3
-----

Online Quarantined Devices across the fabric
-----
Port Index |  PID  |  PWWN
-----
      17   | 051100 | 30:10:00:05:33:ac:c6:13
      17   | 051101 | 30:10:01:05:33:ac:c6:13
-----
```

SEE ALSO [mapsConfig](#), [mapsRule](#), [switchShow](#)

secActiveSize

Displays the size of the active security database.

Synopsis `secactivesize`

Description Use this command to display the size of the active security database. The command also displays the maximum database size.

For switches running Fabric OS v6.2.0 and later, the maximum security database size is 1 megabyte per logical switch. With up to eight partitions, the total database size on a chassis can be up to 8 megabytes. On switches that are not Virtual Fabric-capable, the security database is limited to 1 megabyte. For switches running earlier versions of Fabric OS (up to v5.3.0), the maximum size is 256 Kilobytes.

Notes The effective security DB size is the lowest supported by the fabric. The presence of a Standby CP that runs an earlier version of the operating system will drop the effective security DB size on an Active CP that runs Fabric OS v6.2.0.

The Brocade 200E is unable to handle the maximum DB size supported in v6.2.0 and issues a compact flash warning when the active security database is close to the 1 MB limit.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the size of the active security database:

```
switch:admin> secactivesize
  Size of security active data: 35 bytes \
  (Max 1048576 bytes)
```

SEE ALSO [secDefineSize](#), [secGlobalShow](#)

secAuthSecret

Manages the DH-CHAP shared secret key information.

Synopsis `secauthsecret --show`
 `secauthsecret --set`
 `secauthsecret --remove value | --all`

Description Use this command to manage the DH-CHAP shared secret key database used for authentication. This command displays, sets, and removes shared secret key information from the database or deletes the entire database. If you are performing set or remove operations, when the command is completed new data is saved persistently. New data is effective with the next authentication request. The configuration applies to a switch instance only.

Port level authentication security must be enabled before encryption configuration can be enabled. Pre-shared secret keys should be configured on both ends of the ISL to perform authentication. For encrypted ports, an authentication key of 32 characters is recommended. Spaces are not allowed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--show Lists the WWNs for which a shared secret is configured. In Access Gateway mode, since you cannot configure using domain ID or switchname, these fields will be displayed as -1 and Unknown respectively.

--set Sets shared secrets. You can set shared secrets for F_Port, N_Port in Access Gateway mode and E, EX_Port and F_Port on Fabric OS. This command is interactive. In Access Gateway mode, you can specify only a WWN. In Fabric OS, you can specify a WWN or switchname or domain ID. Spaces are not allowed.

--remove [*wwn* | *domain* | *swname*] Removes the specified WWN entry from the database. If a domain name is specified, it is converted to a WWN and then the entry is removed. If no option is specified, the command is interactive. In Access Gateway mode, you can specify only a WWN. In Fabric OS, you can specify a WWN or switchname or domain ID.

--remove --all Deletes the entire secret key database.

Examples To list the shared secret WWN:

```
switch:admin> secauthsecret --show

WWN                               DIId   Name
-----
10:00:00:60:69:80:5b:e8           1      switch
```

To list the shared secret WWN in Access Gateway mode:

```
switch:admin> secauthsecret --show

WWN                               DId   Name
-----
10:00:00:60:69:80:5b:e8          -1    Unknown
```

To set the shared secret:

```
switch:admin> secAuthSecret --set
```

This command sets up secret keys for the DH-CHAP authentication. The minimum length of a secret key is 8 characters and maximum 40 characters. Setting up secret keys does not initiate DH-CHAP authentication. If switch is configured to do DH-CHAP, it is performed whenever a port or a switch is enabled.

Warning: Please use a secure channel for setting secrets. Using an insecure channel is not safe and may compromise secrets.

Following inputs should be specified for each entry.

1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.

Press Enter to start setting up shared secrets >

```
Enter WWN, Domain, or switch name (Leave blank when done):
10:00:00:60:69:80:05:14
Enter peer secret:
Re-enter peer secret:
Enter local secret:
Re-enter local secret:
```

```
Enter WWN, Domain, or switch name (Leave blank when done):
Are you done? (yes, y, no, n): [no] y
Saving data to key store... Done.
```

To delete the entire secret key database:

```
switch:admin> secAuthSecret --remove --all
```

This command deletes database of DH-CHAP secret keys. If a fabric requires authentication, deleting this database may cause switch to segment from the fabric.

```
Do want to remove secret key database?
(yes, y, no, n): [no] y
Deleting secret key database... Done.
```

SEE ALSO None

secCertUtil

Manages certificates on a switch.

Synopsis

seccertutil

seccertutil genkey [-nowarn] [-keysize 1024 | 2048 | 4096 | 8192]

seccertutil delkey [-nowarn][-all]

seccertutil genscr [-hash sha1 | sha256] [-country *country code*]
 [-state *state*] [-locality *locality*] [-org *organization*]
 [-orgunit *organization unit*] [-cn *common name*]

seccertutil delcsr [-nowarn]

seccertutil showcsr

seccertutil generate [-fcap] [-keysize 1024 | 2048]
 [-hash sha1 | sha256] [-nowarn]

seccertutil generate [-commoncert] [-keysize 1024 | 2048]
 [-hash sha1 | sha256] [-nowarn]

seccertutil generate -https [-keysize 1024 | 2048 | 4096 | 8192] [-type rsa | dsa]
 [-hash sha1 | sha256 | sha512] [-years *value*] [-nowarn]

seccertutil delete [-ldapcert [*certificate name*]] | -syslogcert *certificate_name* |
 -fcapcert | -fcapswcert | -fcapall |
 -commoncertall | *certificate name*] [-nowarn]

seccertutil export [-ldapcert [-certname *certificate name*]] |
 -syslogcert [-certname *certificate_name*] |
 -fcapswcert | -fcapswcsr | -fcapcert |
 -commonswcert | -commonswcsr | -commoncert]
 [-protocol ftp | scp] [-ipaddr *IP address*]
 [-remotedir *remote directory*] [-login *login name*]
 [-password *password*]

seccertutil import [-ldapcert | -syslogcert | -fcapswcert |
 -fcapcert | -config cacert | -config swcert [-enable https]
 -commonswcert | -commoncert]
 [-protocol ftp | scp] [-ipaddr *IP address*]
 [-remotedir *remote directory*] [-certname *certificate name*]
 [-login *login name*] [-password *password*]

seccertutil show [-ldapcert] | [-syslogcert [*certificate name*]] | [-fcapall] |
 [-fcapswcert] | [-commoncertall] | [*filename*]

Description Use this command to manage third-party certificates on a switch, including Public Key Infrastructure (PKI) based certificates, Lightweight Directory Access Protocol (LDAP) certificates, FCAP certificates, and syslog CA certificates. This command also imports or exports Certificate Signing Requests (CSRs) from or to a remote host. This command supports IPV4 and IPV6 addresses.

Use this command to do the following:

- Generate a public/private key pair.
- Delete a public/private key pair.
- Generate a CSR.
- Delete a CSR.
- List existing certificates on a switch.
- Display the contents of a certificate or CSR.
- Delete a specified certificate.
- Import or export a certificate.
- Configure a SSL certificate file name.
- Enable secure protocols.

This command takes an action and associated arguments. If only an action is specified, this command prompts interactively for input values of the associated arguments. The command runs noninteractively when the arguments associated with a given action are specified on the command line. When invoked without operands, this command displays the usage.

This command is also supported in Access Gateway mode, for FCAP authentication between AG and the switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Before you import a certificate from Windows system, convert the certificate to a Unix file format with the **dos2unix** utility.

Operands This command has the following operands:

genkey

Generates a public/private key pair. This is the first step in setting up a third-party certificate. The following operands are optional; when omitted, the command prompts interactively for input values to these operands.

-keysize 1024 | 2048 | 4096 | 8192

Specifies the size of the key. Valid values are 1024, 2048, 4096, or 8192 bits. The greater the value, the more secure is the connection; however, performance degrades with size. The keys are generated only after all existing CSRs and certificates have been deleted.

-nowarn

Specifies that no warning is given when overwriting or deleting data. If this operand is omitted, the command prompts for confirmation before existing CSRs and certificates are deleted.

- delkey**
- Deletes all public/private key pairs with the exception of Encryption-related certificates and key pairs. This command prompts for confirmation unless **-nowarn** is specified.
- nowarn**
- Executes the delete operation without confirmation.
- all**
- Deletes all public/private key pairs including Encryption-related certificates and key pairs.
- gencsr**
- Generates a new CSR for the switch. This is the second step in setting up a third-party certificate. The following operands are optional; if omitted, the command prompts for answers to a series of questions. If only one or a few operands are specified, the command prompts for input to the remaining questions. When all questions are answered, a CSR is generated and placed in a file named *ip_address.csr*, where *ip_address* is the IP address of the switch.
- hash sha1 | sha256**
- Sets the hash type. Valid values are "sha256" and "sha1".
- country *country code***
- Specifies the country. Provide a two-letter country code, for example, US.
- state *state***
- Specifies the state. Provide the full name, for example, California. If the state consists of multiple words, it must be enclosed in double quotes.
- locality *locality***
- Specifies the city. Provide the full name, for example, "San Jose". If the locality consists of multiple words, it must be enclosed in double quotes.
- org *organization***
- Specifies the organization. Provide the full name, for example, Brocade. If the organization consists of multiple words, it must be enclosed in double quotes.
- orgunit *organization unit***
- Specifies the organizational unit, for example, your department name. If the organizational unit consists of multiple words, it must be enclosed in double quotes.
- cn *common name***
- Specifies the common name. Provide a fully qualified Domain Name, or IP address.
- showcsr**
- Displays the content of the CSR on the switch without page breaks. Use the pipe operator followed by the "more" option to display the content one page at a time.
- delcsr [-nowarn]**
- Deletes the CSR. This command prompts for confirmation before deleting the CSR unless **-nowarn** is specified.
- generate**
- Generates the certificate.
- fcap**
- Generates FCAP switch CSR, passphrase, and key.

- commoncert**
Generates the CSR and key file with key size of 1024 or 2048 for the common certificate.
- keysize 1024 | 2048**
Specifies the size of the key. Valid values are 1024 or 2048. bits. The greater the value, the more secure is the connection; however, performance degrades with size.
- hash sha1 | sha256**
Specifies the hash type. Valid values are "sha256" or "sha1".
- nowarn**
Generates the certificate without confirmation.
- https**
Generates HTTPS certificate.
 - keysize 1024 | 2048 | 4096 | 8192**
Specifies the size of the key. Valid values are 1024, 2048, 4096, or 8192. bits. The greater the value, the more secure is the connection; however, performance degrades with size.
 - type rsa | dsa**
Specifies the key pair that the certificate is based on (RSA or DSA).
 - hash sha1 | sha256 | sha512**
Specifies the hash type. Valid values are "sha512", "sha256", or "sha1".
 - years *value***
Specifies the number of years the certificate is valid for.
- delete**
Deletes the specified certificate. Specify one of the following operands:
 - ldapcert**
Deletes all LDAP CA certificates. You may specify a file name when deleting the LDAP CA certificates. The file name is optional. All LDAP certificates are concatenated in a single file. Use **show -ldapcert** to list that file name on the switch. You cannot delete an individual certificate.
 - syslogcert *certificate_name***
Deletes the specified syslog CA certificate. You may specify **-all** to delete all the syslog CA certificates. The syslog certificates are stored in separate files. Use **show -syslogcert** to list the imported CA certificate names.
 - fcapcert**
Deletes the FCAP CA certificate.
 - fcapswcert**
Deletes the FCAP SW certificate.
 - fcapall**
Deletes all FCAP objects including the switch CSR, keys, pass phrase, Root CA certificate, and switch certificate.
 - commoncertall**
Deletes the CSR, key file with key size of 1024 or 2048, switch certificates for the common certificate.

certificate name

The name of the certificate to be deleted.

-nowarn

Deletes the specified file without confirmation. This operand is optional.

export

Exports a CSR to a host. This command is typically used to submit a CSR to the Certification Authority (CA) that issues the certificate. The following operands are optional; if omitted, the command prompts interactively for your input.

-ldapcacert -certname *certificate name*

Exports an LDAP CA certificate from the switch to a remote host. You must specify the name of the LDAP CA certificate to be exported. Use the **show** option for a list of existing certificates.

-syslogcacert [-certname *certificate_name*]

Exports a syslog CA certificate from the switch to a remote host. You must specify the name of the syslog CA certificate to be exported. Use the **show -syslogcacert** option for the list of existing certificates.

-fcapswcert

Exports the FCAP switch certificate from the switch to the remote host. The certificate is exported with the same filename it had when it was imported.

-fcapswcsr

Exports the FCAP switch CSR certificate from the switch to the remote host. The FCAP switch CSR is exported as switchIPfcapSw.csr.

-fcapcacert

Exports the FCAP CA certificate from the switch to the remote host. The FCAP CA certificate is exported as switchIPfcapRootCA.pem.

-commonswcert

Exports the switch certificate from the switch to the remote host. It is exported with the switch IP address as the file name.

-commonswcsr

Exports the switch CSR certificate from the switch to the remote host. The certificate name is in the format is *ip_address*commonSw.csr.

-commoncacert

Exports the CA certificate from the switch to the remote host. The certificate name is in the format is *ip_address*commonRootCA.pem.

-protocol ftp | scp

Specifies the protocol as either FTP or SCP.

-ipaddr *IP address*

Specifies the IP address of the remote host.

-remotedir *remote directory*

Specifies the remote directory. Provide a fully qualified path name.

-login *login name*

Specifies the login name for the server.

-password *password*

Specifies the password for the user account. When using SCP, for security reasons, do not enter a password on the command line. Use the interactive version instead. Use **CTRL-C** to skip the password.

import

Imports a certificate. Use this command to import a certificate from the server or to download a certificate issued by a CA after sending the CSR to the CA. If the certificate type is not mentioned, a generic certificate is imported. Specify one of the following certificates:

-ldapcert

Imports an LDAP CA certificate.

-syslogcert

Imports a syslog CA certificate.

-fcapswcert

Imports an FCAP switch certificate.

-fcapcert

Imports an FCAP CA certificate.

-config cacert

Imports a CA certificate.

-config swcert [-nowarn]

Imports an SSL certificate. If **-nowarn** is specified, the SSL certificate import is blocked without warning if the certificate is invalid. If **-nowarn** is omitted, you are prompted for confirmation before importing an invalid switch certificate. However, you will be able to import the invalid certificate. The **-nowarn** option effectively prevents you from importing an invalid swcert certificate.

-enable https

Enables secure https. Optionally use this operand with **-config swcert** only.

The following operands are optional with **import**; if omitted, the command prompts interactively for your input.

-commonswcert

Imports the switch certificate. Optionally, you can enable HTTPS using the **-enable https** and **-config swcert** options.

-commoncacert

Imports the CA certificate.

-protocol ftp | scp

Specifies the protocol as either FTP or SCP.

-ipaddr *IP address*

Specifies the IP address of the remote host.

-remotedir *remote directory*

Specifies the remote directory. Provide a fully qualified path name.

-certname *certificate name*

Specifies the certificate name.

- login** *login name*
Specifies the login name for the server.
- password** *password*
Specifies the password for the user account. When using SCP, for security reasons, do not enter a password on the command line. Use the interactive version instead.
- show**
Lists all existing PKI-based certificates on the switch. The following operands are optional and exclusive.
- ldapcert**
Displays the LDAP certificate file name.
- syslogcert** [*certificate name*]
Displays the list of imported syslog CA certificates. You may specify a file name to display content of a particular CA certificate.
- fcapall**
Indicates whether FCAP-related files exist or are empty.
- fcapswcert**
Displays the contents of the FCAP SW certificate.
- commoncertall**
Indicates whether CommonCert-related files exist or are empty.
- filename**
Displays the contents of the specified certificate.

Examples To generate a public/private key pair in interactive mode:

```
switch:admin> seccertutil genkey

Generating a new key pair will automatically do the following:
1. Delete all existing CSRs.
2. Delete all existing certificates.
3. Reset the certificate filename to none.
4. Disable secure protocols.

Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
```

To generate a public/private key pair in noninteractive mode:

```
switch:admin> seccertutil genkey -nowarn -keysize 1024
Generating new rsa public/private key pair
Done.
```

To generate a CSR in interactive mode:

```
switch:admin> seccertutil gencsr
Input hash type (sha1 or sha256): sha1
Country Name (2 letter code, eg, US):US
State or Province Name (full name, \
eg, California):California
```

```

Locality Name (eg, city name):"San Jose"
Organization Name (eg, company name):Brocade
Organizational Unit Name (eg, department \
or section name): IT
Common Name (Fully qualified Domain Name, \
or IP address):192.168.38.206
generating CSR, file name is: 192.168.38.206.csr
Done

```

To generate a CSR in noninteractive mode:

```

switch:admin> seccertutil genscr -country US \
-state California -locality "San Jose" \
-org Brocade -orgunit software -cn 192.168.38.206
Input hash type (sha1 or sha256): sha1
generating CSR, file name is: 192.168.38.206.csr
Done

```

To delete the CSR in interactive mode:

```

switch:admin> seccertutil delcsr
WARNING!!!

About to delete CSR: /etc/fabos/certs/sw0/10.20.5.82.csr
ARE YOU SURE (yes, y, no, n): [no] y

```

To delete a CSR in noninteractive mode:

```

switch:admin> seccertutil delcsr -nowarn

```

To import an LDAP certificate from a remote host to the local switch in interactive mode:

```

switch:admin> seccertutil import -ldpcacert
Select protocol [ftp or scp]: ftp
Enter IP address: 195.168.38.206
Enter remote directory: /users/home/remote_certs
Enter certificate name (must have ".cert", \
".cer" or ".pem" suffix): ldap.cer
Enter Login Name: mylogin
Enter Password: password
Success: imported certificate [ldap.cert].

```

To import an LDAP certificate from a remote host to the local switch in noninteractive mode:

```

switch:admin> seccertutil import -ldpcacert \
-protocol ftp -ipaddr 195.168.38.206 \
-remotedir /users/home/remote_certs -certname ldap.cer \
-login abcd -passwd passwd
Success: imported certificate [ldap.cert].

```

To import an FCAP switch certificate in interactive mode:

```
switch:admin> seccertutil import -fcapswcert
Select protocol [ftp or scp]: scp
Enter IP address: 10.32.2.25
Enter remote directory: /users/myname/OPENSSL
Enter certificate name (must have ".pem" suffix):O1.pem
Enter Login Name: myname
myname@10.32.2.25's password:
Success: imported certificate [O1.pem].
```

To import an FCAP CA certificate in interactive mode:

```
switch:admin> seccertutil import -fcapcacert
Select protocol [ftp or scp]: scp
Enter IP address: 10.32.2.25
Enter remote directory: /users/myname/OPENSSL
Enter certificate name (must have ".pem" suffix):CACert.pem
Enter Login Name: myname
myname@10.32.2.25's password:
Success: imported certificate [CACert.pem].
```

To import a PKI-based certificate with configure and enable option in interactive mode:

```
switch:admin> seccertutil import -config swcert -enable https
Select protocol [ftp or scp]: ftp
Enter IP address: 192.168.38.206
Enter remote directory: /users/home/mycerts
Enter certificate name (must have ".crt" or ".cer" \
".pem" or "psk" suffix): filename.crt
Enter Login Name: username
Enter Password: password
Success: imported certificate [filename.crt].
Certificate file in configuration has been updated.
Secure http has been enabled.
```

To import a PKI-based certificate with configure and enable option in noninteractive mode:

```
switch:admin> seccertutil import -config swcert \
-enable https -protocol ftp -ipaddr 195.168.38.206 \
-remotedir /users/home/remote_certs -certname file.crt \
-login abcd -password passwd
Success: imported certificate [file.crt].
Certificate file in configuration has been updated.
Secure http has been enabled.
```

To export the switch certificate from the switch to the remote host in interactive mode:

```
switch:admin> seccertutil import -commonswcert
Select protocol [ftp or scp]: scp
Enter IP address: 10.32.2.25
Enter remote directory: /users/myname/OPENSSL
Enter certificate name (must have ".pem" suffix):O1.pem
Enter Login Name: myname
myname@10.32.2.25's password:
Success: imported certificate [O1.pem].
```

To export the switch certificate from the switch to the remote host in noninteractive mode:

```
switch:admin> seccertutil import -commonswcert \
  -ipaddr 10.32.2.25 -remotedir /users/myname/OPENSSL \
  certs -certname 01.pem -login myname -password passwd
Success: imported certificate [01.pem].
```

To export an LDAP CA certificate from the local switch to a remote host in interactive mode:

```
switch:admin> seccertutil export -ldapcert
Select protocol [ftp or scp]: scp
Enter IP address: 192.168.38.206
Enter remote directory: /users/home/remote/mycerts
Enter Login Name: username
Enter LDAP certificate name (must have ".pem" \
  or ".crt" or ".cer" suffix): ldap.cer
Password: password
Success: exported LDAP certificate
```

To export an LDAP CA certificate from the local switch to a remote host in noninteractive mode:

```
switch:admin> seccertutil export -ldapcert \
  -protocol ftp -ipaddr 192.168.38.206 \
  -remotedir /users/home/remote_certs \
  -login abcd -passwd passwd -certname ldap.cer
Success: exported LDAP certificate
```

To export an FCAP switch certificate in interactive mode:

```
switch:admin> seccertutil export -fcapswcert
Select protocol [ftp or scp]: scp
Enter IP address: 10.32.2.25
Enter remote directory: /users/myname/OPENSSL
Enter Login Name: myname
myname@10.32.2.25's password:
Success: exported FCAP switch certificate
```

To export an FCAP switch CSR in interactive mode:

```
switch:admin> seccertutil export -fcapswcsr
Select protocol [ftp or scp]: scp
Enter IP address: 10.32.2.25
Enter remote directory: /users/myname/OPENSSL
Enter Login Name: myname
myname@10.32.2.25's password:
Success: exported FCAP switch CSR
```

To delete an LDAP CA certificate in interactive mode:

```
switch:admin> seccertutil delete -ldapcert filename.pem
WARNING!!!
About to delete certificate: filename.cer
ARE YOU SURE (yes, y, no, n): [no] y
Deleted LDAP certificate successfully
```

To delete an LDAP CA certificate in noninteractive mode:

```
switch:admin> seccertutil delete -ldapcert filename.pem
Deleted LDAP certificate successfully
```

To delete all FCAP PKI objects:

```
switch:admin> seccertutil delete -fcapall
WARNING!!!

About to delete FCAP-Files:
ARE YOU SURE (yes, y, no, n): [no] y
```

To generate FCAP PKI objects:

```
switch:admin> seccertutil generate \
-fcapall -keysize 1024
Generating a new key pair will automatically
do the following:
1. Delete all existing CSRs.
2. Delete existing switch certificate.
3. Reset the certificate filename to none.

Continue (yes, y, no, n): [no] y
Installing Private Key and Csr...
Switch key pair and CSR generated...
```

To display the FCAP PKI objects:

```
switch:admin> seccertutil show -fcapall
List of fcap files:
Passphrase      : Exist
Private Key     : Exist
CSR             : Exist
Certificate     : Empty
Root Certificate: Exist
```

To display the LDAP certificate file name on a switch:

```
switch:admin> seccertutil show -ldapacert
List of ldap ca certificate files:
Sample.cer
```

To display the content of the CSR one page at a time:

```
switch:admin> seccertutil showcsr | more
```

To display the contents of an FCAP SW certificate:

```
switch:admin> seccertutil show -fcapswcert
```

To display the contents of a PKI-based or LDAP CA certificate:

```
switch:admin> seccertutil show filename
```

SEE ALSO None

secCryptoCfg

Configures and displays cryptographic parameters and templates.

Synopsis

```
seccryptocfg --replace -type SSH | https [-cipher cipher_string |
-kex value | -mac value] -force

seccryptocfg --default -type https | SSH [-force]

seccryptocfg --apply template_name

seccryptocfg --import template_name [-server ip_address -name user
-proto scp | sftp | ftp -file remote_file_name]

seccryptocfg --export template_name [-server ip_address -name user
-proto scp | sftp | ftp -file remote_file_name]

seccryptocfg --ltemplates

seccryptocfg --verify template_name

seccryptocfg --delete template_name

seccryptocfg --show [template_name]

seccryptocfg --help
```

Description Use this command to perform the following functions:

- Configure cryptographic parameters such as ciphers, key exchange (kex) algorithm, and message authentication code (MAC) algorithm for SSH or HTTPS.
- Reset the cryptographic parameters to default.
- Set the template configuration.
- Import a template file.
- Export a template file.
- Display the list of templates available.
- Verify the configuration against the template.
- Delete a template file.
- Display the configured cryptographic or template parameters.

When the cryptographic parameters are configured, this command restarts the SSH or HTTPS daemon for the changes to take effect. This command prompts for confirmation unless you use the **-force** option.

A maximum of eight templates including the default templates are supported. You cannot overwrite the default configurations but can upload the configurations, edit, and then download it with a different name. You can create a new template similar to default templates, download, and apply. For more information on the format and rule of the template, refer to *Brocade Fabric OS Administrator's Guide*.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--replace -type SSH | https

Configures the specified ciphers, key exchange, and MAC algorithms for SSH or HTTPS and restarts the SSH or HTTPS daemon.

-cipher *cipher_string*

Specifies the SSH or HTTPS cipher list. You can either specify one or more ciphers separated by a comma, for example, 3des-cbc,aes128-cbc,aes192-cbc. The following ciphers are configured by default. You can modify the cipher list using this operand.

For SSH:

- aes128-ctr
- aes192-ctr
- aes256-ctr
- aes128-cbc
- 3des-cbc
- aes192-cbc
- aes256-cbc

For HTTPS:

- ECDH
- DH
- HIGH
- MD5
- CAMELLIA
- SRP
- PSK
- AESGCM

- kex *value***
 Specifies the SSH kex algorithms list. You can either specify one or more kex algorithms separated by a comma, for example, ecdh-sha2-nistp256,diffie-hellman-group-exchange-sha256. The following kex algorithms are configured by default. You can modify the kex algorithms list using this operand.
- ecdh-sha2-nistp256
 - ecdh-sha2-nistp384
 - ecdh-sha2-nistp521
 - diffie-hellman-group-exchange-sha256
 - diffie-hellman-group-exchange-sha1
 - diffie-hellman-group14-sha1
 - diffie-hellman-group1-sha1
- mac *value***
 Specifies the SSH MAC list. You can either specify one or more MAC algorithms separated by a comma, for example, hmac-md5,hmac-sha1. The following MAC algorithms are configured by default. You can modify the MAC algorithms list using this operand.
- hmac-md5
 - hmac-sha1
 - hmac-sha2-256
 - hmac-sha2-512
- default -type SSH | https**
 Resets the cipher, kex, and MAC configurations to default.
- force**
 Executes without prompting for a confirmation.
- apply *template_name***
 Sets a default or user-defined template file. The following default templates are supported:
- default_generic
 - default_strong
 - default_fips
 - default_cc
- import**
 Imports a template file from a specified external host.
- export**
 Exports a template file to the specified external host.
- template_name***
 Specifies the template name.
- server *ip_address***
 Specifies the IP address of the remote host.
- name *user***
 Specifies the user name for the host. Depending on your host configuration, this command may prompt for a password.

- proto scp | sftp | ftp**
Specifies the protocol as either SCP, SFTP, or FTP.
- file remote_file_name**
Specifies the remote directory. Provide a fully qualified path name.
- lstemplates**
Displays the list of templates available.
- verify template_name**
Verifies the running configuration against a required configuration specified in the template file.
- delete template_name**
Deletes the specified template file.
- show [template_name]**
Displays the configured cipher, kex, and MAC algorithms. If the template name is specified, this command displays the configuration contents in the specified template file.
- help**
Displays the command usage.

Examples

To configure ciphers, kex, and MAC for SSH and HTTPS:

```
switch:admin> seccryptocfg --replace -type SSH -cipher \
    3des-cbc,aes128-cbc,aes192-cbc -kex diffie-hellman-group-exchange-sha1 \
    -mac hmac-sha2-256
```

To configure ciphers for SSH:

```
switch:admin> seccryptocfg --replace -type SSH -cipher \
    aes128-ctr,aes192-ctr
This command requires the daemon(s) SSH to be restarted.
Existing sessions will be terminated.
Please confirm and provide the preferred option
Press Yes(Y,y), No(N,n) [N]:y
```

To configure ciphers for HTTPS:

```
switch:admin> seccryptocfg --replace -type https -cipher \
    '!ECDH:!DH:HIGH:-MD5:!CAMELLIA'
This command requires the daemon(s) HTTP to be restarted.
Existing sessions will be terminated.
Please confirm and provide the preferred option
Press Yes(Y,y), No(N,n) [N]:no
```

To configure ciphers for HTTPS using the **-force** option:

```
switch:admin> seccryptocfg --replace -type https -cipher \
    '!ECDH:!DH:HIGH:-MD5:!CAMELLIA' -force
HTTP cipher list configured successfully.
```

To set the default CC configuration template:

```
switch:admin> seccryptocfg --apply default_cc
Validating....
Applying...

Template configurations applied successfully
```

To import a template file:

```
switch:admin> seccryptocfg --import cc_import -server 10.70.12.10 \
-name brocade -proto scp -file /users/home40/brocade/default_cc
brocade@10.70.12.10's password:

Import successful
```

To export a template file:

```
switch:admin> seccryptocfg --export default_cc -server 10.70.12.10 \
-name brocade -proto scp -file /users/home40/brocade
brocade@10.70.12.10's password:

Export successful
```

To display the list of templates available:

```
switch:admin> seccryptocfg --ltemplates

List of templates:
default_generic
default_fips
cc_import
default_strong
default_cc
```

To verify a template file:

```
switch:admin> seccryptocfg --verify default_strong
Validating ....
Verifying ....
Failed for ....
    SSH:Kex,Mac,Enc

Failed: System configuration is not compliant with input template
switch:admin> seccryptocfg --verify default_cc
Validating ....
Verifying ....
Passed: System configuration is compliant with the input template
```

To delete a template file:

```
switch:admin> seccryptocfg --delete cc_import
```

To display the configured ciphers, kex, and MAC algorithms:

```
switch:admin> seccryptocfg --show
HTTPS Cipher List      : !ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM
SSH Cipher List       : aes128-ctr,aes192-ctr,aes256-ctr,aes128-cbc, \
                        3des-cbc,aes192-cbc,aes256-cbc
SSH Kex Algorithms List : ecdh-sha2-nistp256,ecdh-sha2-nistp384, \
                        ecdh-sha2-nistp521, \
                        diffie-hellman-group-exchange-sha256, \
                        diffie-hellman-group-exchange-sha1, \
                        diffie-hellman-group14-sha1, \
                        diffie-hellman-group1-sha1
SSH MACs List         : hmac-md5,hmac-sha1,hmac-sha2-256,hmac-sha2-512
```

To display configuration contents in a specific template file:

```
switch:admin> seccryptocfg --show default_cc
[Ver] 0.1
[SSH]
Enc:aes128-cbc,aes256-cbc
Kex:diffie-hellman-group14-sha1,ecdh-sha2-nistp256,ecdh-sha2-nistp384,\
    ecdh-sha2-nistp521
Mac:hmac-sha1,hmac-sha2-256,hmac-sha2-512
[AAA]
RAD_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!SSLv3
LDAP_Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!SSLv3
[HTTPS]
Ciphers:!ECDH:!DH:HIGH:-MD5:!CAMELLIA:!SRP:!PSK:!AESGCM:!3DES:!DES
```

SEE ALSO None

secDefineSize

Displays the size of the defined security database.

Synopsis `secdefinesize`

Description Use this command to display the size of the defined security database. The command also displays the maximum database size.

For switches running Fabric OS v6.2.0 and later, the maximum security database size is 1 Megabyte per logical switch. With up to eight partitions, the total database size on a chassis can be up to 8 Megabytes. On switches that are not Virtual Fabric-capable, the security database is limited to 1 Megabyte. For switches running earlier versions of Fabric OS (up to v5.3.0), the maximum size is 256 Kilobytes.

The effective security DB size is the lowest supported by the fabric. The presence of a Standby CP that runs an earlier version of the operating system drops the effective security DB size on an Active CP that runs Fabric OS v6.2.0 or later.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the size of the defined security database

```
switch:admin> secdefinesize
  Size of security defined data: 35 bytes \
  (Max 1048576 bytes)
```

SEE ALSO [secActiveSize](#), [secGlobalShow](#)

secGlobalShow

Displays the current internal security state information.

Synopsis `secglobalshow`

Description Use this command to display security server-specific information as a snapshot of its current state. The output may include information about the following:

- General security parameters
- The latest zone transaction
- The current status of the RCS transaction
- The state of the Domain
- wwnDL state

This command is intended primarily for debugging purposes by technical support staff. The information displayed may not be supported between releases and is subject to change without notice.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To view the current security state:

```
switch:admin> secglobalshow

----General Security Information----
flag 1, saveflag 0
transId 0
Queue Size 0
final Rca 0
reliablemsg 0
reliablePhase 0
Primary pub key: Empty
Primary Version 0
Primary WWN 10:00:00:05:1e:01:23:e0 (local switch)
Stamp 0
----The latest zone transaction--
last retVal from zone: not used
last zone size when calling zone: not used
----The latest RCS STATUS----
RCS was enabled
sec_aca: free
RCS latest Phase: Completion
RCS Message ==> RCS transaction completes.
----DataBase STATUS----
Retry Role 0
Retry Query 0
Security Active DataSize 35 bytes

-----Domain State-----
```

secGlobalShow

```
Active Sum 215b
Security Defined DataSize 35 bytes
Define Sum 215b
Zone Size (include enabled configuration) 312 bytes
Zone sum e04b215b
sec_db: free
primaryDLPhase 0
----wnnDL State-----
pid      tid      key      sec      usec

----- LOG CACHE -----
14:08:50 813905136 secipadm_ipchange receives notification
14:08:50 850588912 secProcessIPChange starts processing
14:08:50 850588912 secProcessIPChange acks completion
(Output truncated)
```

SEE ALSO [secActiveSize](#), [secDefineSize](#)

secHelp

Displays information about security commands.

Synopsis **sechelp**

Description Use this command to display a list of security commands with a brief description of the commands.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display a list of security commands:

```
switch:admin> sechelp
fipscfg          Configure FIPS mode operation
secactivesize    Displays size of the active (security)
                 database
secauthcertificate Configure FCAP switch certificate properties
secauthsecret    Creates/Manages DHCHAP secret key details
seccertutil      Creates/Manages/Displays third party PKI
                 certificates
secdefinesize    Displays size of the defined (security)
                 database
secglobalshow    Displays current internal security state
                 information
secpolicyabort   Aborts changes to defined policy
secpolicyactivate Activates all policy sets
secpolicyadd     Adds members to an existing policy
secpolicycreate  Creates a new policy
secpolicydelete  Deletes an existing policy
secpolicydump    Displays all members of existing policies
secpolicyfcsmove Moves a member in the FCS policy
secpolicyremove  Removes members from an existing policy
secpolicysave   Saves defined policy set and sends to all
                 switches
secpolicyshow    Shows members of one or more policies
secstatsreset    Resets security statistics
secstatsshows    Displays security statistics
sshutil          Configure SSH authentication options
```

SEE ALSO None

secPolicyAbort

Aborts all unsaved changes to the defined database.

Synopsis `secpolicyabort`

Description Use this command to abort all changes to the defined security database that have not been saved to flash memory and to abort changes to policy creation and modification operations from all the switches if a fabric-wide consistency policy is not set for the particular policy.

Notes When an FCS policy is enabled, this command can be issued only from the Primary FCS switch.

Only the user who made the changes to the defined database may use this command to abort them.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To abort all changes that have not been saved to nonvolatile memory:

```
primaryfcs:admin> secpolicyabort
Unsaved data has been aborted.
primaryfcs:admin> secpolicyabort
No new data to abort.
```

SEE ALSO [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#), [secStatsShow](#)

secPolicyActivate

Saves and activates the Defined Security Policy Set.

Synopsis **secpolicyactivate**

Description Use this command to activate the current defined security policy to all switches in the fabric. This activates the policy set on the local switch or all switches in the fabric depending on the fabric-wide consistency policy.

If there are changes to the SCC, DCC, or FCS policies in the current CLI or API transaction that have not been saved to the Defined Security Policy Set, then this command saves the changes to the Defined Security Policy Set first, and then activates it. If there are no changes, but the Defined Security Policy Set differs from the Active Security Policy Set, then the Defined Security Policy Set is activated. If there are no changes and the Defined Security Policy Set is the same as the Active Security Policy Set, then nothing is done.

After activation the defined policy set becomes the Active Policy Set.

Use **secPolicyShow** to display the members of an existing policy in the Active or Defined Security Policy Sets.

Notes The behavior of this command is the same for tolerant and strict fabric-wide consistency.

When an FCS policy is enabled, this command can be issued only from the Primary FCS switch.

Any modifications to the SCC, DCC, and FCC DB are saved and activated. When **secPolicyActivate** is issued after the **secPolicySave** command, it might fail.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To activate the defined security policy set on all switches in the fabric:

```
switch:admin> secpolicyactivate
About to overwrite the current Active data.
ARE YOU SURE (yes, y, no, n): [no] y
secpolicyactivate command was completed successfully.
```

SEE ALSO [fddCfg](#), [secPolicyAbort](#), [secPolicyAdd](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#), [secPolicyShow](#)

secPolicyAdd

Adds members to an existing security policy.

Synopsis `secpolicyadd "name","member[;member...]" [-legacy]`

Description Use this command to add one or more members to an existing access policy.

Each policy corresponds to a management method. The list of members of a policy acts as an access control list for that management method. Before a policy is created, there is no enforcement for that management method; all access is granted. After a policy has been created and a member has been added to the policy, that policy becomes closed to all access except from included members. If all members are then deleted from the policy, all access is denied for that management method (the DCC_POLICY is an exception).

Attempting to add a member to a policy that already is a member causes this command to fail.

In a Virtual Fabric Environment, when you create a DCC lockdown policy on a logical switch, the DCC policy is created for each port in the chassis, even though the ports are not currently present in the local logical switch. This is done to provision the DCC policy for the ports that may be moved later. If a policy seems stale at any point, use **secPolicyDelete** to remove all stale DCC policies.

Fabric-wide consistency policies can be configured on per logical switch basis, which applies the FCS policy to the corresponding fabric connecting to the logical switch. Automatic policy distribution for DCC, SCC and FCS remains unchanged in Fabric OS v6.2.0 and can be configured on a per logical switch basis.

On switches running Fabric OS v7.1.0 or later, all DCC and SCC security policy members are sorted based on their world wide names (WWNs) in order to avoid a segmentation of ports. This is not the case for switches running earlier firmware versions; on these switches, security member lists are unsorted. When a switch with an unsorted security policy member list tries to join a switch that runs Fabric OS v7.1.0 or later and is configured with an ordered security policy list, port segmentation occurs because of mismatching security policy lists. To prevent this from happening, use the **-legacy** option to add security policy members in a manner that matches the order of security policy members in Fabric OS v7.0.0 and earlier.

Notes When an FCS policy is enabled, this command can be issued only from the Primary FCS switch. The **secpolicyadd** command can be issued on all switches for SCC and DCC policies as long as fabric-wide consistency policy is not set for the particular policy.

Do not add the WWNs of front or translate (xlate) domains to the FCS policy if the edge fabric is connected to an FC Router.

Backup FCS switches typically cannot modify the policy. However, if the Primary FCS switch in the policy list is not reachable, then a backup FCS switch is allowed to modify the policy. If all the reachable backup FCS switches are running pre-v5.3.0 versions of Fabric OS, a non-FCS v5.3.0 switch is allowed to modify the policy so that a new switch can be added to the policy.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

"name"

Specify the name of an existing policy to which you want to add members. Valid values for this operand include the following:

- DCC_POLICY_ *nnn*
- FCS_POLICY
- SCC_POLICY

The specified policy name must be capitalized.

The DCC_POLICY_ *nnn* name has the common prefix DCC_POLICY_ followed by a string of user-defined characters. These characters do not have to be capitalized like regular policy names, but they are case-sensitive.

"member"

Specify a list of one or more member switches to be included in the security policy. The list must be enclosed in quotation marks; members must be separated by semicolons. Depending on the policy type, members are specified as follows.

FCS_POLICY or SCC_POLICY Members

This policy type requires member IDs to be specified as WWN strings, Domains, or switch names. If Domain IDs or switch names are used, the switches associated must be present in the fabric or the command fails.

DCC_POLICY Members

The DCC_POLICY_ *nnn* is a list of device port names associated with a specific switch and port index combination. An empty DCC_POLICY does not stop access to the switch. The device port name is specified by its port WWN string. The switch and port index combination must be in the *switch port* format, where *switch* can be specified as a WWN, a domain, or a switch name, and *port* is specified by port numbers separated by commas and enclosed in either brackets or parentheses; for example, (2, 4, 6). Ports enclosed in brackets include the devices currently attached to those ports.

The following examples illustrate several ways to specify the port values:

(1-6)

Selects ports 1 through 6.

(*)

Selects all ports on the switch.

[3, 9]

Selects ports 3 and 9 and all devices attached to those ports.

[1-3, 5]

Selects ports 1 through 3 and 5 and all devices attached to those ports.

[*]

Selects all ports on the switch and devices currently attached to those ports.

-legacy

Adds a security policy with the desired order.

Examples To add a member to the SCC_POLICY using the device WWN:

```
primaryfcs:admin> secpolicyadd "SCC_POLICY", \
    "12:24:45:10:0a:67:00:40"
Member(s) have been added to SCC_POLICY.
```

To add two devices to attach to domain 3, ports 1 and 3, in an existing empty DCC policy; the port WWN of the first device is 11:22:33:44:55:66:77:aa and port WWN of the second device is 11:22:33:44:55:66:77:bb:

```
primaryfcs:admin> secpolicyadd "DCC_POLICY_abc", \
    "11:22:33:44:55:66:77:aa;11:22:33:44:55:66:77:bb;3(1,3)"
Member(s) have been added to DCC_POLICY_abc.
```

To add a security policy:

```
switch:admin> secpolicyadd SCC_POLICY,\
    "10:00:00:05:1e:a3:01:d9" -legacy
switch:admin> secpolicyshow
```

```
-----
ACTIVE POLICY SET
-----
DEFINED POLICY SET
SCC_POLICY
  WWN                               Did  swName
-----
10:00:00:05:1e:a1:ef:b9             2   sw0
10:00:00:05:1e:a3:00:59             5   sw0
10:00:00:05:1e:a2:f9:09             3   sw0
10:00:00:05:1e:a3:01:d9             3   sw0
```

SEE ALSO [distribute](#), [fddCfg](#), [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#)

secPolicyCreate

Creates a new security policy.

Synopsis `secpolicycreate "name" [, "member[;member...]" [-legacy]`

Description Use this command to create a new policy and to edit Switch Connection Control (SCC), Device Connection Control (DCC), and Fabric Configuration Server (FCS) policies on the local switch. All policies can be created only once, except for the DCC_POLICY_ *nnn*. Each DCC_POLICY_ *nnn* must have a unique *name*. This command can be issued on all switches in the current fabric for SCC and DCC policies if they are not intended to be fabric-wide.

Adding members while creating a policy is optional. You can add members to a policy later, using the **secPolicyAdd** command.

Each policy corresponds to a management method. The list of members of a policy acts as an access control list for that management method. Before a policy is created, there is no enforcement for that management method, which is all access is granted. After a policy is created and a member is added to the policy, that policy is closed to all access except to included members. If all members are then deleted from the policy, all access is denied for that management access method.

All newly created policies are saved on the local switch only, unless the switch has a fabric-wide consistency policy for that policy.

In a Virtual Fabric environment, when you create a DCC lockdown policy on a logical switch, the DCC policy is created for each port in the chassis, even though the ports are not currently present in the local logical switch. This is done to provision the DCC policy for the ports that may be moved later. If a policy seems stale at any point, use the **secPolicyDelete** command to remove all stale DCC policies.

Fabric wide consistency policies can be configured on a logical switch basis, which applies the FCS policy to the corresponding fabric connecting to the logical switch. Automatic policy distribution behavior for DCC, SCC and FCS remains unchanged in Fabric OS v6.2.0 or later and can be configured on a logical switch basis.

On switches running Fabric OS v7.1.0 or later, all DCC and SCC security policy members are sorted based on their world wide names (WWNs) in order to avoid a segmentation of ports. This is not the case for switches running earlier firmware versions; on these switches, security member lists are unsorted. When a switch with an unsorted security policy member list tries to join a switch that runs Fabric OS v7.1.0 or later and is configured with an ordered security policy list, port segmentation occurs because of mismatching security policy lists. To prevent this from happening, use the **-legacy** option to create security policy members in a manner that matches the order of security policy members in Fabric OS v7.0.0 and earlier.

Notes When an FCS policy is enabled, this command can be issued only from the Primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

"name"

Specify the name of the policy you want to create. Valid values for this operand include the following:

- DCC_POLICY_ *nnn*
- SCC_POLICY
- FCS_POLICY

The specified policy name must be capitalized.

The DCC_POLICY_ *nnn* name has the common prefix DCC_POLICY_ followed by a string of user-defined characters. These characters do not have to be capitalized like regular policy names. Valid values for DCC_POLICY_ *nnn* are user-defined alphanumeric or underscore characters. The maximum length is 30 characters, including the prefix DCC_POLICY_.

secpolicycreate DCC_POLICY "" may be used to indicate DCC lockdown. This command creates a unique policy for each port in the fabric locking it down to the device connected or creating an empty policy to disallow any device to be connected to it. This can be done only when there are no other DCC policies defined on the switch. The switch must be in an enabled state for DCC lockdown to succeed. On a disabled switch, executing **secpolicycreate** DCC_POLICY "" will not create any DCC policies.

"member"

Specify one or more members to be included in the security policy. The member list must be enclosed in double quotation marks and members separated by semicolons. The member list must be separated from the name field by a comma and a space. Depending on the policy type, members are specified as follows:

DCC_POLICY Members

The DCC_Policy_ *nnn* is a list of devices associated with a specific switch and port index combination. An empty DCC_POLICY does not stop access to the switch. The device is specified by its port WWN. The switch and port combination must be in the switch *port* format.

switch can be specified using a WWN, domain, or switch name.

port can be specified by port numbers separated by commas and enclosed in either brackets or parentheses: for example, (2, 4, 6). Ports enclosed in brackets include the devices currently attached to those ports.

The following examples illustrate several ways to specify the port values:

(1-6)

Selects ports 1 through 6.

(*)

Selects all ports on the switch.

[3, 9]

Selects ports 3 and 9 and all devices attached to those ports.

[1-3, 5]

Selects ports 1 through 3 and 5 and all devices attached to those ports.

[*]

Selects all ports on the switch and devices currently attached to those ports.

SCC_POLICY and FCC_POLICY Members

This policy type requires member IDs to be specified as WWN strings, domains, or switch names. If domain or switch names are used, the switches associated must be present in the fabric or the command fails.

To add all switches in the current fabric as members of the policy, enter an asterisk enclosed in quotation marks (*) as the member value. This feature cannot be used by the other security commands.

-legacy

Creates a security policy with desired order.

Examples

To create an FCS policy (While creating the FCS policy, the local switch WWN is automatically included in the list. Switches included in the FCS list are FCS switches and the remaining switches in the fabric are non-FCS switches. Out of the FCS list, the switch that is in the first position becomes the Primary FCS switch and the remaining switches become backup FCS switches. If the first switch in the FCS list is not reachable, the next switch becomes the Primary):

```
primaryfcs:admin> secpolicycreate "FCS_POLICY", "3; 4"
FCS_POLICY has been created.
```

To create a device policy to allow two devices to attach to domain 3 ports 1 and 3 (the WWN of first device is 11:22:33:44:55:66:77:aa and the WWN of second device is 11:22:33:44:55:66:77:bb):

```
primaryfcs:admin> secpolicycreate "DCC_POLICY_aB_7", \
"11:22:33:44:55:66:77:aa;11:22:33:44:55:66:77:bb;3[1,3]"
DCC_POLICY_abc has been created.
```

To create a SCC policy in a fabric with three switches:

1. Check if a policy exists.

```
switch:admin> secpolicyshow
```

```
-----
ACTIVE POLICY SET
-----
DEFINED POLICY SET
```

2. Identify switches in the fabric.

```
switch:admin> fabricshow
```

```
Switch ID      Worldwide Name      Enet IP Addr FC IP Addr Name
-----
2:fffc02 10:00:00:05:1e:39:5f:67 10.32.69.53 10.20.30.53 "sw1"
                fec0:60:69bc:60:260:69ff:fe80:d4a
4:fffc04 10:00:00:05:1e:04:ef:0e 10.32.69.49 10.20.30.49 'sw2"
                fec0:60:69bc:54:205:1eff:fe04:ef0e
200: fffc08 10:00:00:05:1e:35:cd:ef 10.32.69.117 10.20.30.117 \
                "nSW4100_98"
```

3. Create a SCC policy that includes switches with domain IDs 2 and 4.

```
switch:admin> secpolicycreate "SCC_POLICY","2;4"
SCC_POLICY has been created.
```

4. Activate the policy.

```
switch:admin> secpolicyactivate
About to overwrite the current Active Policy Set.
ARE YOU SURE (yes, y, no, n): [no] y
secpolicyactivate command was completed successfully.
```

To create an SCC policy that includes all switches in the fabric:

```
switch:admin> secpolicycreate "SCC_POLICY",*
SCC_POLICY has been created.
```

To create a security policy:

```
switch:admin> secpolicycreate SCC_POLICY,\
"10:00:00:05:1e:a1:ef:b9; 10:00:00:05:1e:a3:00:59; \
10:00:00:05:1e:a2:f9:09" -legacy
SCC_POLICY created as user specified.
switch:admin> secpolicyshow
```

```
-----
ACTIVE POLICY SET
-----
DEFINED POLICY SET
SCC_POLICY
WVN                                DId swName
-----
10:00:00:05:1e:a1:ef:b9           2 sw0
10:00:00:05:1e:a3:00:59           5 sw0
10:00:00:05:1e:a2:f9:09           3 sw0
```

SEE ALSO

[fddCfg](#), [distribute](#), [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#), [secStatsShow](#)

secPolicyDelete

Deletes an existing security policy.

Synopsis `secpolicydelete name`

Description Use this command to delete an existing security policy from the defined security database. Run **secPolicyActivate** to delete the policies from the active security policy list. Deleting a security policy does not cause any traffic disruption.

Each policy corresponds to a management method. The list of members of a policy acts as an access control list for that management method. Before a policy is created, there is no enforcement for that management method; all access is granted. After a policy has been created and a member has been added to the policy, that policy becomes closed to all access except from included members. If the policy is deleted all access is granted.

Notes When an FCS policy is enabled, this command can be issued only from the Primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operand is required:

"*name*"

Specify the name of a security policy to delete. The policy name must be capitalized. Quotation marks are optional. Once a security policy is deleted, fabric-wide switch access through that method is unrestricted. Valid security policy names include the following:

DCC_POLICY_nnn

Deletes the specified Device Connection Control (DCC) policy. The DCC_POLICY_nnn name has the common prefix DCC_POLICY_ followed by a string of user-defined characters. These characters do not have to be capitalized.

ALL_DCC_POLICY

Deletes all DCC policies from the defined policy list.

ALL_STALE_DCC_POLICY

Deletes all stale DCC policies from the defined policy list. DCC policies become stale when the ports are removed from a logical switch.

SCC_POLICY

Deletes the Switch Connection Control policy from the defined policy list.

FCS_POLICY

Deletes the Fabric Configuration Server policy from the defined policy list.

Examples To delete an existing security policy:

```
switch:admin> secpolicydelete "DCC_POLICY_ab1"
About to delete policy DCC_POLICY_ab1.
Are you sure (yes, y, no, n):[no] y
DCC_POLICY has been deleted.
```

secPolicyDelete

To delete all existing DCC policies in the fabric:

```
primaryfcs:admin> secpolicydelete ALL_DCC_POLICY
About to clear all the DCC policies
ARE YOU SURE (yes, y, no, n): [no] y
```

To delete all stale DCC policies in the fabric:

```
primaryfcs:admin> secpolicydelete ALL_STALE_DCC_POLICY
About to clear all STALE DCC policies
ARE YOU SURE (yes, y, no, n): [no] y
```

SEE ALSO

[secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyCreate](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#), [secPolicyShow](#)

secPolicyDump

Displays the members of one or all existing security policies.

Synopsis `secpolicydump ["listtype"],"name"]`

Description Use this command to display, without page breaks, the members of an existing policy in the active and defined (saved) databases. When issued without operands, this command displays the members of all security policies.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command can be issued from all FCS switches in the fabric.

Operands This command has the following optional operands:

"listtype"

Specifies the database to display, enclosed in double quotation marks. The name for an active database is "Active"; the name for a saved, defined database is "Defined". If *listtype* is not specified, all databases are displayed.

"name"

Specifies the security policy for which to display the members. Valid values for this operand include the following:

- DCC_POLICY_*nnn*
- FCS_POLICY
- SCC_POLICY

The specified policy name must be capitalized and enclosed in double quotation marks.

The DCC_POLICY_*nnn* name has the common prefix DCC_POLICY_ followed by a string of user-defined characters. These characters do not have to be capitalize. If *name* is not specified, all existing policies are displayed.

Examples To display all security policy information from all databases without page breaks:

```
switch:admin> secpolicydump
ACTIVE POLICY SET
FCS_POLICY
Pos Primary WWN                               DId swName
-----
1   Yes   10:00:00:60:69:30:15:5c 1 primaryfcs
2   No    10:00:00:60:69:30:1e:62 4 switch
-----
DEFINED POLICY SET
FCS_POLICY
Pos Primary WWN                               DId swName
-----
1   Yes   10:00:00:60:69:30:15:5c 1 primaryfcs
2   No    10:00:00:60:69:30:1e:62 4 switch
-----
```

To display all security policies in the active database:

```
switch:admin> secpolicydump "active"
```

```

-----
ACTIVE POLICY SET
FCS_POLICY
Pos   Primary WWN                               DID swName
-----
  1   Yes    10:00:00:05:1e:39:5f:67 3 NeptuneSec
  2   No     10:00:00:05:1e:90:09:4a - Unknown

SCC_POLICY
WWN                               DID swName
-----
10:00:00:05:1e:39:5f:67  3 NeptuneSec
10:00:00:05:1e:90:09:4a  - Unknown

DCC_POLICY_h1
Type   WWN                               DID swName
-----
Switch 10:00:00:05:1e:39:5f:67  3 NeptuneSec.
=Index=> 34.
Device 21:00:00:e0:8b:13:5e:8d
Device 21:00:00:e0:8b:13:5e:8e

```

To display all security policies in the defined database:

```
switch:admin> secpolicydump "Defined"
```

```

-----
DEFINED POLICY SET
FCS_POLICY
Pos   Primary WWN                               DID swName
-----
  1   Yes    10:00:00:05:1e:39:5f:67 3 NeptuneSec
  2   No     10:00:00:05:1e:90:09:4a - Unknown

SCC_POLICY
WWN                               DID swName
-----
10:00:00:05:1e:39:5f:67  3 NeptuneSec
10:00:00:05:1e:90:09:4a  - Unknown

DCC_POLICY_h1
Type   WWN                               DID swName
-----
Switch 10:00:00:05:1e:39:5f:67  3 NeptuneSec.
=Index=> 34.
Device 21:00:00:e0:8b:13:5e:8d
Device 21:00:00:e0:8b:13:5e:8e

```

To display the FCS policies in the defined database:

```
switch:admin> secpolicydump "Defined","FCS_POLICY"
```

```

-----
DEFINED POLICY SET
FCS_POLICY
Pos   Primary WWN                               DId swName
-----
  1   Yes    10:00:00:05:1e:39:5f:67 3 NeptuneSec
  2   No     10:00:00:05:1e:90:09:4a - Unknown

```

To display the SCC policies in the defined database:

```
switch:admin> secpolicydump "Defined","SCC_POLICY"
```

```

-----
DEFINED POLICY SET
SCC_POLICY
WWN                               DId swName
-----
10:00:00:05:1e:39:5f:67 3 NeptuneSec
10:00:00:05:1e:90:09:4a - Unknown

```

To display the SCC policies in the active database:

```
switch:admin> secpolicydump "Active","SCC_POLICY"
```

```

-----
ACTIVE POLICY SET
SCC_POLICY
WWN                               DId swName
-----
10:00:00:05:1e:39:5f:67 3 NeptuneSec
10:00:00:05:1e:90:09:4a - Unknown

```

SEE ALSO [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyCreate](#), [secPolicyDelete](#), [secPolicyRemove](#), [secPolicySave](#), [secPolicyShow](#)

secPolicyFCSMove

Moves a member in the FCS policy.

Synopsis `secpolicyfcsmove [from, to]`

Description Use this command to move an FCS member from one position to another position in the FCS list. Only one FCS can be moved at a time. The first FCS switch in the list that is also present in the fabric is the Primary FCS.

Notes If a backup FCS is moved to the first position, it becomes the primary FCS after activation.

An FCS policy must be enabled to execute this command, and the command must be issued from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

from

Specify the position of the FCS switch you want to move.

to

Specify the position to which you want to move the FCS switch.

Examples To move the backup FCS switch from position 2 to position 3 in the FCS list (interactively):

```
switch:admin> secpolicyfcsmove
Pos   Primary WWN                               DId swName.
=====
  1   Yes      10:00:00:60:69:10:02:18   1 switch5.
  2   No       10:00:00:60:69:00:00:5a   2 switch60.
  3   No       10:00:00:60:69:00:00:13   3 switch73.

Please enter position you'd like to move \
from : (1..3) [1] 2

Please enter position you'd like to move \
to   : (1..3) [1] 3

-----
                DEFINED POLICY SET
-----
FCS_POLICY
Pos   Primary WWN                               DId swName
-----
  1   Yes      10:00:00:60:69:10:02:18   1 switch5.
  2   No       10:00:00:60:69:00:00:13   3 switch73.
  3   No       10:00:00:60:69:00:00:5a   2 switch60.
-----
```


To move Backup FCS switch from position 3 to position 1 in the FCS list(non-interactively):

```
switch:admin> secpolicyshow
```

```

ACTIVE POLICY SET
1  Yes      10:00:00:05:1e:39:5f:67  2  switch1
2  No       10:00:00:05:1e:04:ef:0e  4  switch2
3  No       10:00:00:05:1e:35:cd:ef 200 switch3

```

```
switch:admin> secpolicyfcsmove 3,1
```

```

DEFINED POLICY SET
FCS_POLICY
Pos  Primary WWN                      DId swName
-----
1  No       10:00:00:05:1e:35:cd:ef 200 switch3
2  Yes      10:00:00:05:1e:39:5f:67  2  switch1
3  No       10:00:00:05:1e:04:ef:0e  4  switch2

```

```
switch:admin> secpolicyactivate
```

```

About to overwrite the current Active Policy Set.
ARE YOU SURE (yes, y, no, n): [no] y
secpolicyactivate command was completed successfully.

```

```
switch:admin> secpolicyshow
```

```

ACTIVE POLICY SET
FCS_POLICY
Pos  Primary WWN                      DId swName
-----
1  Yes      10:00:00:05:1e:35:cd:ef 200 switch3
2  No       10:00:00:05:1e:39:5f:67  2  switch1
3  No       10:00:00:05:1e:04:ef:0e  4  switch2

```

SEE ALSO [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyCreate](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#), [secPolicyShow](#)

secPolicyRemove

Removes members from an existing security policy.

Synopsis `secpolicyremove "name" ,"member[;member...]"`

Description Use this command to remove one or more members from an existing security policy. It is not possible to remove all members from the FCS_POLICY; the local switch WWN cannot be deleted from the FCS policy. In the case of SCC policy, if it is empty after removing all members, all access to the switch itself would be disallowed.

Beginning Fabric OS v7.3, this command will not remove the local WWN from SCC policy list in HIF and non-HIF mode. The command also prompts for confirmation before removing any member if HIF is enabled.

Notes If an FCS policy is enabled, this command must be issued from the primary FCS switch.

After removing members from an existing security policy, execute the **secPolicyActivate** command to activate the current defined policy.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

"name"

Specify the name of an existing policy you want to remove members from. This operand is required. Valid values for this operand include the following:

- DCC_POLICY_ *nnn*
- FCS_POLICY
- SCC_POLICY

The specified policy name must be capitalized.

The DCC_POLICY policy name has the common prefix DCC_POLICY_ followed by a string of user-defined characters. These characters do not have to be capitalized like regular policy names, but are case-sensitive.

"member"

Specify a member or list of members to delete from the policy. The list must be enclosed in quotation marks; members must be separated by semicolons. This operand is required.

Depending on the policy type, members can be specified using IP address, WWN, domain, or switch name.

WWN Member Policy Types

The following policy types require members be specified by WWN address:

- FCS_POLICY
- SCC_POLICY

These policy types require member IDs be specified as WWN strings, domains, or switch names. If domain or switch names are used, the switches associated must be present in the fabric or the command fails.

DCC_POLICY Members

The DCC_Policy_nnn is a list of devices associated with a specific switch and port combination. The device is specified with a WWN string. The switch and port combination must be specified in the *switch port* format where *switch* can be specified by switch WWN, domain, or switch name. The *port* parameter can be specified by port number separated by commas, and enclosed in either brackets or parentheses: for example, (2, 4, 6). Ports enclosed in brackets include the devices currently attached to those ports.

The following examples illustrate several ways to specify the port values:

| | |
|----------|--|
| (1-6) | Selects ports 1 through 6. |
| (*) | Selects all ports on the switch. |
| [3, 9] | Selects ports 3 and 9 and all devices attached to those ports. |
| [1-3, 5] | Selects ports 1 through 3 and 5 and all devices attached to those ports. |
| [*] | Selects all ports on the switch and devices currently attached to those ports. |

Examples To remove a member that has a WWN of 12:24:45:10:0a:67:00:40 from SCC policy:

```
switch:admin> secpolicyremove "SCC_POLICY", \
"12:24:45:10:0a:67:00:40"
Member(s) have been removed from SCC_POLICY.
```

To remove a member with SCC policy and HIF enabled:

```
switch:admin> secpolicyremove "SCC_POLICY" "10:00:00:27:f8:d0:b3:8f"
HIF mode is enabled. About to remove SCC_POLICY member(s).
ARE YOU SURE (yes, y, no, n): [no] y
```

SEE ALSO [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyCreate](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicySave](#), [secStatsShow](#)

secPolicySave

Saves a defined security policy to persistent memory.

Synopsis **secpolicysave**

Description Use this command to save a defined security policy to persistent memory. **SecPolicySave** saves the modified SCC, DCC, and FCS policies to the Defined Security Policy Set on the local switch.

Notes This command is always a local switch operation. A fabric-wide consistency configuration does not affect the behavior of this command.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To save a defined policy set to persistent memory:

```
switch:admin> secpolicysave
secpolicysave command was completed successfully.
```

SEE ALSO [fddCfg](#), [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyCreate](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secStatsShow](#)

secPolicyShow

Displays an existing security policy including the FCS policy.

Synopsis `secpolicyshow ["policy_set"["", name"]]`

Description Use this command to display the members of an existing policy in the Active or Defined security policy set. The command can be issued from all FCS switches.

This command displays the policy database one page at a time. Use the **secPolicyDump** command to display the policy database without page breaks.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

"policy_type"

Specify which policy to display, in quotation marks. Valid values are "Active", "Defined", or an asterisk (*) for both Active and Defined. This operand is optional. If not specified, all databases are displayed.

"name"

Specify the name of the security policy you want to view, in quotation marks. The specified policy name must be capitalized. This operand is optional. Valid values for this operand include the following:

- DCC_POLICY_ *nnn*
- FCS_POLICY
- SCC_POLICY

The DCC_POLICY_ *nnn* name has the common prefix DCC_POLICY_ followed by a string of user-defined characters. These characters do not have to be capitalized like regular policy names, but they are case-sensitive.

Examples To display all security policies from active databases:

```
switch:admin> secpolicyshow "active"
```

```

ACTIVE POLICY SE
FCS_POLICY

Pos Primary WWN                               DId swName
-----
 1 Yes 10:00:00:60:69:30:15:5c 1 primaryfcs
 2 No  10:00:00:60:69:30:1e:62 4 switch

```

secPolicyShow

To display all security policies from defined databases:

```
switch:admin> secpolicyshow "defined"
```

| DEFINED POLICY | | | | |
|----------------|---------|-------------------------|-----|------------|
| FCS_POLICY | | | | |
| Pos | Primary | WWN | Did | swName |
| 1 | Yes | 10:00:00:60:69:30:15:5c | 1 | primaryfcs |
| 2 | No | 10:00:00:60:69:30:1e:62 | 4 | switch |

SEE ALSO [fddCfg](#), [secPolicyAbort](#), [secPolicyActivate](#), [secPolicyAdd](#), [secPolicyCreate](#), [secPolicyDelete](#), [secPolicyDump](#), [secPolicyRemove](#), [secPolicySave](#)

secStatsReset

Resets one or all security statistics to 0.

Synopsis `secstatsreset [name][,"domain[:domain"]]`

Description Use this command to reset one or all security statistics to 0. This command can be issued on any switch to reset the security statistics on the local switch or chassis. If an FCS policy is enabled and **secStatsReset** is issued on the primary FCS switch, this command can reset security statistics for any or all switches in the fabric.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands When invoked without operands, this command displays the security statistics on the local switch or chassis. The following operands are optional:

name

Specify the name of a security statistic you would like to reset. The specified policy name must be capitalized. If executed on the primary FCS, specify an asterisk (*) to reset all security policies. Valid values for this operand include the following:

- TELNET_POLICY
- HTTP_POLICY
- SCC_POLICY
- DCC_POLICY
- LOGIN
- INVALID_CERT
- AUTH FAIL
- TS_OUT_SYNC
- NO_FCS
- INCOMP_DB
- ILLEGAL_CMD

To access DCC policies, enter DCC_POLICY. Violations are not tracked for individual DCC policies. The statistics for all DCC_POLICY violations are grouped together.

domain(s)

Specify a list of domain IDs on which to reset the security statistics. Specify an asterisk (*) to represent all switches in the fabric or specify a list of domains, separated by semicolons and enclosed in quotation marks. This option can only be executed when an FCS policy is enabled and when the command is issued from the primary FCS switch. When domain is specified, the name operand is required.

Examples To reset all statistics on the local switch:

```
switch:admin> secstatsreset
About to reset all security counters.
ARE YOU SURE (yes, y, no, n):[no] y
Security statistics reset to zero.
```

secStatsReset

To reset DCC_POLICY statistics on domains 1 and 69:

```
primaryfcs:admin> secstatsreset DCC_POLICY, "1;69"  
Reset DCC_POLICY statistic.
```

SEE ALSO [secStatsShow](#)

secStatsShow

Displays one or all security statistics.

Synopsis `secstatshow [name[, "domain[:domain]"]]`

Description Use this command to display one or all security statistics. This command can be issued on any switch to display local security statistics. If an FCS policy is enabled and **secStatsShow** is issued on the primary FCS switch, this command can retrieve and display the security statistics for any or all switches in the fabric.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands When invoked without operands, this command displays the security statistics on the local switch or chassis. The following operands are optional:

name

Specify the name of the security statistic you want to view. The specified policy name must be capitalized. If executed on the primary FCS, specify an asterisk (*) to represent all security policies in the fabric. Valid values for this operand include the following:

- TELNET_POLICY
- HTTP_POLICY
- SCC_POLICY
- DCC_POLICY
- LOGIN
- AUTH FAIL
- TS_OUT_SYNC
- NO_FCS
- INCOMP_DB
- ILLEGAL_CMD

To access DCC policies, enter DCC_POLICY. Violations are not tracked for individual DCC policies. The statistics for all DCC_POLICY violations are grouped together.

The INVALID_CERT policy name is not supported.

domain

Specify one or more domains for which to display the security statistics. Specify an asterisk (*) in quotation marks to represent all switches in the fabric or specify a list of domains separated by semicolons. This option can only be executed when an FCS policy is enabled and the command is issued from the primary FCS switch. When ***domain*** is specified, the ***name*** operand is required.

Examples To display the LOGIN policy statistics for the local domain.

```
switch:admin> secstatsshow LOGIN

Fabric Statistics:

Domain 1:
Name          Value
=====
LOGIN         2
```

To display statistic information for TELNET_POLICY for all switches in the fabric from the primary FCS switch.

```
primaryfsc:admin> secstatsshow TELNET_POLICY,""

Fabric Statistics:

Domain 1:
Name          Value
=====
TELNET_POLICY 0

Domain 69:
Name          Value
=====
TELNET_POLICY 0

Domain 70:
Name          Value
=====
TELNET_POLICY 0
```

SEE ALSO [secStatsReset](#)

sensorShow

Displays sensor readings.

Synopsis `sensorshow`

Description Use this command to display the current temperature, fan, and power supply status and readings from sensors located on the switch. The actual location of the sensors varies, depending on the switch type.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To view all sensor values:

```
switch:admin> sensorshow
sensor 1: (Temperature) is Ok, value is 39 C
sensor 2: (Temperature) is Absent
sensor 3: (Temperature) is Absent
sensor 4: (Temperature) is Absent
sensor 5: (Temperature) is Ok, value is 26 C
sensor 6: (Temperature) is Ok, value is 27 C
sensor 7: (Fan          ) is Ok, speed is 2537 RPM
sensor 8: (Fan          ) is Ok, speed is 2537 RPM
sensor 9: (Fan          ) is Ok, speed is 2556 RPM
sensor 10: (Power Supply ) is Ok
sensor 11: (Power Supply ) is Absent
sensor 12: (Power Supply ) is Ok
sensor 13: (Power Supply ) is Absent
```

SEE ALSO [fanShow](#), [tempShow](#)

serDesTuneMode

Configures and displays SerDes tuning values.

Synopsis

serdestunemode --set

serdestunemode --reset

serdestunemode --show

serdestunemode --autoenable

serdestunemode --autodisable

serdestunemode --autoreset

serdestunemode --autoshow

serdestunemode --help

Description

Use this command to configure and display SerDes tuning values on the Brocade FC8-16 blade and to control tuning optimization of backend links for all Brocade FC8-xx blades.

When used with the **--set**, **--reset** and **--show** options, this command enables or disables the SerDes tuning mode for the Brocade FC8-16 port blade and to display the configuration.

Enabling SerDes tuning mode for the Brocade FC8-16 port blade is indicated under the following conditions:

- You have a Brocade FC8-16 blade installed in slot 2 or in slot 7 (or in both slots) in a DCX-4S chassis with a CR4S-8 core blade.
- You notice increasing numbers of CRC errors on ports 1/11, 2/0, 2/1, 2/3, 2/4, 2/5, 2/8, 2/10, 6/58 and 7/1.

Enabling SerDes tuning mode change the SerDes values on the above-mentioned ports in slots 1, 2 and 7 along with their peer ports, provided the port blades are FC8-16s and the peer is a CR4S-8 core blade. SerDes tuning mode changes take effect immediately and are saved persistently.

When used with the **--autoenable**, **--autodisable**, **--autoreset** and **--autoshow** options, this command configures SerDes tuning optimization for backend SerDes errors. SerDes tuning optimization is indicated when you receive RASLog errors of type C2-5825 indicating a CRC error with a good EOF. Brocade highly recommends that you contact your service provider if you encounter such errors. If the error appears to be an isolated incident, you may enable auto SerDes tuning for a short period of time (24-48 hours) to see if this resolves the problem. Then disable the function to reduce the risk of resetting register values unnecessarily if single occurrence errors are encountered. If the errors persist contact your service provider.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|----------------------|--|
| --set | Enables SerDes tuning mode and sets new SerDes values on FC8-16 blades in slots 1, 2 and 7 and on peer ports in the CR4S-8 core blade. |
| --reset | Disables SerDes tuning mode and resets SerDes values to default on FC8-16 blades in slots 1, 2 and 7 and on peer ports in the CR4S-8 core blade. |
| --show | Indicates whether the SerDes tuning mode is enabled or disabled. |
| --autoenable | Enables SerDes tuning optimization on all backend ports. |
| --autodisable | Disables further tuning optimization but retains all current tuning values identified by the algorithm. |
| --autoreset | Disables further tuning optimization and resets all tuning values back to their default values. If SerDes tune mode or was set, the configuration is restored. |
| --autoshow | Displays whether SerDes tuning optimization is enabled or disabled. |
| --help | Displays the command usage. |

Examples To enable SerDes tuning mode on a DCX-4S chassis with two FC8-16 blades and to display the configuration:

```
switch:admin> serdestunemode --set
serdestunemode is turned on
New serdes values set on slot = 1
New serdes values set on slot = 2
New serdes values set on slot = 7

switch:admin> serdestunemode --show
serdestunemode is enabled
```

To disable SerDes tuning mode and reset the SerDes values to defaults and to display the configuration:

```
switch:admin> serdestunemode --reset
serdestunemode is turned off
New serdes values reset on slot = 1
New serdes values reset on slot = 2
New serdes values reset on slot = 7

switch:admin> serdestunemode --show
serdestunemode is disabled
```

serDesTuneMode

To enable SerDes tuning optimization on all backend ports and to display the configuration:

```
switch:admin> serdestunemode --autoenable  
Autotunemode is turned on
```

```
switch:admin> serdestunemode --show  
Autotunemode is turned on
```

To disable SerDes tuning optimization on all backend ports and retain all current tuning values:

```
switch:admin> serdestunemode --autodisable  
Autotunemode is turned off
```

```
switch:admin> serdestunemode --show  
Autotunemode is turned off
```

To disable SerDes tuning optimization on all backend ports and reset all tuning values to their default values:

```
switch:admin> serdestunemode --autoreset  
AutoTune Resetting Default Tuning On Next Poll
```

```
switch:admin> serdestunemode --show  
Autotunemode is in reset state
```

SEE ALSO None

setContext

Sets the logical switch context to a specified FID.

Synopsis `setcontext [FID | switchname]`

Description Use this command to set the logical switch context to a specified fabric ID (FID) or unique switch name. The FID uniquely defines a partition as a logical switch. Use `lscfg --show` to display currently configured partitions and their FIDs.

A logical switch context defines the boundaries within which a user can execute commands in a Virtual Fabric-aware environment. In a Virtual Fabric-aware environment, all commands are context-specific. When a user executes a switch-wide command, the command applies to the current logical switch context.

On legacy platforms, or if a logical switch context is not set explicitly, switch commands apply to the default logical switch context (FID 128). When the context is changed, switch-wide commands apply to the new logical switch context.

You must have chassis permissions to access all logical switches in the chassis and to use the `setContext` command to change the current context to any partition configured on the chassis. A user without chassis permissions can change contexts only within the list of FIDs specified in the user's access permissions. Refer to `userConfig` for more information.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

If a logical switch FID is deleted, users logged in to that switch are logged out.

Operands This command has the following operand:

| | |
|-------------------|--|
| <i>FID</i> | Specifies the fabric ID of the logical switch instance for which the context is set. |
| <i>switchname</i> | Specifies the name of the logical switch for which the context is set. |

Examples To change the logical switch context to FID 20:

```
switch:admin> setcontext 20
```

To change the logical switch context to switch_20:

```
switch:admin> setcontext switch_20
```

SEE ALSO [IsCfg](#), [userConfig](#)

setDbg

setDbg

Sets the debug level of the specified module.

Synopsis `setdbg [module_name][level]`

Description Use this command to set the debug level of a specified module. Debug levels filter the display of debug messages to the serial console. By default, no debug messages are displayed.

High debug level values can generate a large volume of messages, degrading the system response time.

The set of supported modules and their current debug levels are displayed by the command **dbgShow**.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

module_name

Specify the name of the module for which you want to set the debug levels. Module names are case-sensitive. If this operand is omitted an error message is displayed.

level

Specify the debug level for the specified module (0 to 9). A zero (0) value (default) specifies that no messages are to display. Higher values cause more messages from that module to display. If this operand is omitted an error message is displayed.

Examples To set the debug level for a module named NS to value 3:

```
switch:admin> setdbg NS 3
switch:admin> dbgshow NS
Module NS,      debug level = 3, verbose level = 0
```

SEE ALSO [dbgShow](#)

setVerbose

Specifies module verbose level.

Synopsis `setverbose [module_name][level]`

Description Use this command to set the verbose level of the specified module. These levels filter the display of the debug message to the serial console. By default, no debug messages are displayed.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

module_name Specify the name of the module for which the verbose level is to be set; module names are case-sensitive.

level Specify the verbose level (0 to 9).

Examples To set the verbose level of a module named NS to value 3:

```
switch:admin> setverbose NS 3
switch:admin> dbgshow NS
Module NS, debug level = 0, verbose level = 3
```

SEE ALSO [dbgShow](#), [setDbg](#)

sfpShow

Displays Small Form-factor Pluggable (SFP) transceiver information.

Synopsis

sfpshow

sfpshow [*slot*]*port* -link [-force]

sfpshow [*slot*]*geport* [-f]

sfpshow -pid *pid*

sfpshow -all

sfpshow -health

sfpshow --help

Description

Use this command to display information about Serial Identification SFPs, also known as module definition "4" SFPs. These SFPs provide extended information that describes the SFP capabilities, interfaces, manufacturer, and other information.

Use this command without operands to display a summary of all SFPs in the switch. For each port, the summary displays the SFP type and, for serial ID SFP, the vendor name and SFP, serial number, and speed capability (in Gbps). Refer to **switchShow** for an explanation of the two-letter codes.

Use this command with a port number to display detailed information about the serial ID SFP in the specified port. In this mode, this command displays values described in the "Gigabit Interface Converter" spec by Sun Microsystems, et al. The detailed view also displays the total number of Monitoring and Alerting Policy Suite (MAPS) state transitions for each SFP and a time stamp of the last polling time. The power-on time is displayed for 16 Gbps/32 Gbps SFPs and QSFPs only.

Use the **-all** operand to display detailed information for all available SFPs. On switches running Fabric OS v7.0.0 or later, this view includes the power on time for the switch in years and hours.

For "smart" SFPs including mini-SFPs, this command displays additional fields, including module temperature, voltage, received optical power, transmitted optical power (long wave only), laser diode drive current, optional status/control register, alarm and warning flags, as well as high and low thresholds programmed on the SFPs.

A "Can not read Serial Data!" message indicates that an SFP ID could not be detected on that location (it could be a fixed port type, or an unsupported SFP type, or an interchassis link without a cable).

Notes

The identifier field has the value **3** for SFPs and mSFPs. mSFPs are supported only on FC8-64 blades and can be identified with OUI 00:05:33.

The **sfpShow** command does not reflect changes in the **sfpShow** output if any SFPs are replaced or removed while a port or a switch is disabled.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The identifier field has the following values to indicate the various transceiver types:

- 1 for GBIC
- 2 for On-board
- 3 for SFP
- 6 for XFP
- 12 for QSFP
- 13 for QSFP+
- 17 for QSFP28

Operands

This command has the following operands:

slot

For bladed systems only, specifies the slot number of the port to display, followed by a slash (/)

port

Specifies the number of the port for which to display the SFP information, relative to its slot for bladed systems. Use **switchShow** for a list of valid ports. This operand is optional; if omitted, this command displays a summary of all SFPs on the switch.

-link [-force]

Displays the diagnostic information from cached data for a local switch port and the peer port. The **-force** option displays the real-time diagnostic information from the registers for a port.

-f

Refreshes the SFP information. This option is valid only when **sfpShow** is issued for a specific port.

-pid *pid*

Specifies the diagnostics information of a remote switch port along with its peer port.

-all

Displays detailed data for all available SFPs on the switch. This operand is not compatible with *slot/port* operands.

-health

Displays SFP health status information from MAPS for 10 Gbps and 16 Gbps SFPs and for 16 Gbps QSFPs (Quad SFPs). When any of the diagnostic SFP parameters such as current, voltage, receiver power, transmit power, and temperature crosses user-configured low and high thresholds, the SFP health state changes. Possible health states include the following:

Green

SFP operates within MAPS thresholds.

Yellow

At least one SFP parameter monitored by MAPS has crossed a low or high threshold.

No License

The switch does not have a Fabric Vision license.

Unknown

This state displays for 8 Gbps, 4 Gbps and 2 Gbps SPF types not supported by this feature.

Paused

Health monitoring is not enabled on the switch.

--help

Displays the command usage.

Examples

To display SFP information including SFP health parameters on a Brocade DCX 8510-8:

```
switch:admin> sfpshow -health
Slot 1/Port 0: id
Slot 1/Port 1: id
Slot 1/Port 2: id
Slot 1/Port 3: id
Slot 1/Port 4: id
Slot 1/Port 5: id
Slot 1/Port 6: id
Slot 1/Port 7: id
Slot 1/Port 8: id
Slot 1/Port 9: id
Slot 1/Port 10: id
Slot 1/Port 11: id
Slot 1/Port 12: id
Slot 1/Port 13: id
Slot 1/Port 14: id
Slot 1/Port 15: id
Slot 2/Port 0: id
Slot 2/Port 1: id
Slot 2/Port 2: id
Slot 2/Port 3: id
Slot 2/Port 4: id
Slot 2/Port 5: id
Slot 2/Port 6: id
Slot 2/Port 7: id
Slot 2/Port 8: id
Slot 2/Port 9: id
Slot 2/Port 10: id
Slot 2/Port 11: id
Slot 2/Port 12: id
Slot 2/Port 13: id
Slot 2/Port 14: id
Slot 2/Port 15: id
Slot 3/Port 0: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000HJ2 Speed: 4,8,16_Gbps Health: Green
Slot 3/Port 1: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000HN2 Speed: 4,8,16_Gbps Health: Green
Slot 3/Port 2: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000HZ2 Speed: 4,8,16_Gbps Health: Green
Slot 3/Port 3: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000HS2 Speed: 4,8,16_Gbps Health: Green
Slot 3/Port 4: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000HM2 Speed: 4,8,16_Gbps Health: Green
Slot 3/Port 5: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000HY2 Speed: 4,8,16_Gbps Health: Green
Slot 3/Port 6: id (sw) Vendor: BROCADE \
    Serial No: HAA110491000J62 Speed: 4,8,16_Gbps Health: Green
Output truncated
```

To display detailed information about a single 16G QSFP:

```
switch:user> sfpshow 3/44
QSFP No: 11 Channel No:0
Identifier:      13   QSFP+
Connector:      12   MPO Parallel Optic
Transceiver:    0000000000000000 16_Gbps  sw Short_dist
Encoding:       5    64B66B
Baud Rate:     140   (units 100 megabaud)
Length 9u:     0    (units km)
Length 50u (OM4): 100 (units m)
Length 62.5u: 0    (units m)
Vendor Name:   BROCADE
Vendor OUI:    00:05:1e
Vendor PN:     57-1000294-01
Vendor Rev:    A
Wavelength:   850   (units nm)
Options:       00000fde
Max Case Temp: 70    (C)
Device Tech:   0x00
Serial No:     HUA1140700000A1
Date Code:    140215
DD Type:      0x8
Enh Options:   0x0
Status/Ctrl:   0x0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature:   32    Centigrade
Current:       7.502 mAmps
Voltage:       3272.6 mVolts
RX Power:     -1.8   dBm (659.9uW)
State transitions: 1
Last poll time: 02-16-2016 PST Tue 22:18:49
```

To display SFP information when a new SFP is inserted in a disabled port:

```
switch:user> sfpshow 0
Identifier: 3   SFP
Connector: 7   LC
Transceiver: 540c404000000000 2,4,8_Gbps M5,M6 sw Short_dist
Encoding: 1    8B10B
Baud Rate: 85   (units 100 megabaud)
Length 9u: 0    (units km)
Length 9u: 0    (units 100 meters)
Length 50u: 5    (units 10 meters)
Length 62.5u:2  (units 10 meters)
Length Cu: 0    (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-1000012-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 003a Loss_of_Sig,Tx_Fault,Tx_Disable
BR Max: 0
BR Min: 0
Serial No: UAF109280000J24
Date Code: 090711
DD Type: 0x68
Enh Options: 0xfa
Status/Ctrl: 0x0
```

sfpShow

```
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature:    Not Available
Current       :    Not Available
Voltage       :    Not Available
RX Power      :    Not Available
TX Power      :    Not Available
Last poll time:    Polling has not started
```

To display detailed information about all SFPs:

```
switch:user> sfpshow -all
(output truncated)
=====
Slot 8/Port 50:
=====

=====
Slot 8/Port 51:
=====

=====
Slot 8/Port 52:
=====
QSFP No: 13 Channel No:0
Identifier: 13 QSFP+
Connector: 12 MPO Parallel Optic
Transceiver: 0000000000000000 16_Gbps id
Encoding: 5 64B66B
Baud Rate: 140 (units 100 megabaud)
Length 9u: 0 (units km)
Length 50u: 25 (units 2 meters)
Length 62.5u:0 (units 1 meter)
Length Cu: 0 (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-0000090-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 00000fde
Max Case Temp: 70 (C)
Device Tech: 0x00
Serial No: HTA110491002833
Date Code: 101207
DD Type: 0x8
Enh Options: 0x0
Status/Ctrl: 0x0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 38 Centigrade
Current: 6.394 mAmps
Voltage: 3266.0 mVolts
RX Power: -0.1 dBm (977.0uW)

State transitions: 1
Last poll time: 02-03-2012 UTC Fri 07:36:05
(output truncatad)
```

To display mini-SFPs on a X6 with a FC32-48 blade:

```
switch:user> setcontext 52
switch:user> sfpshow

Slot 8/Port 34: id (sw) Vendor: BROCADE \
    Serial No: AA1041J0134 Speed: 4,8,16_Gbps
Slot 8/Port 35: --
Slot 8/Port 36: id (sw) Vendor: BROCADE \
    Serial No: HAA21338102SEP2 Speed: 4,8,16_Gbps
Slot 8/Port 37: id (sw) Vendor: BROCADE \
    Serial No: JAF115100000F7T Speed: 8,16,32_Gbps
Slot 8/Port 38: --
Slot 8/Port 39: --
Slot 8/Port 40: id (sw) Vendor: BROCADE \
    Serial No: HAA21403104VFR2 Speed: 4,8,16_Gbps
Slot 8/Port 41: --
Slot 8/Port 42: --
Slot 8/Port 43: id (sw) Vendor: BROCADE \
    Serial No: HAF11102000175A Speed: 4,8,16_Gbps
```

To display mini-SFP details on a X6 with a FC32-48 blade:

```
switch:user> sfpshow 8/43
Identifier: 3 SFP
Connector: 7 LC
Transceiver: 7004404000000000 4,8,16_Gbps \
    M5 sw Short_dist
Encoding: 6 64B66B
Baud Rate: 140 (units 100 megabaud)
Length 9u: 0 (units km)
Length 9u: 0 (units 100 meters)
Length 50u (OM2): 3 (units 10 meters)
Length 50u (OM3): 10 (units 10 meters)
Length 62.5u:0 (units 10 meters)
Length Cu: 0 (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-0000088-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 003a Loss_of_Sig, \
    Tx_Fault,Tx_Disable
BR Max: 0
BR Min: 0
Serial No: HAF11102000175A
Date Code: 110107
DD Type: 0x68
Enh Options: 0xf2
Status/Ctrl: 0xb0
Pwr On Time: 1.08 years (9463 hours)
E-Wrap Control: 0
O-Wrap Control: 0
Alarm flags[0,1] = 0x5, 0x0
Warn Flags[0,1] = 0x5, 0x0
Temperature: 37 Centigrade
Current: 7.298 mAmps
Voltage: 3353.0 mVolts
```

```
RX Power:-2.2 dBm (598.0uW)
TX Power:-2.8 dBm (525.9 uW)
```

```
State transitions: 1
Last poll time: 03-17-2016 UTC Thu 02:45:52
```

To display detailed information for GbE port on a X6 with a SX6 blade:

```
switch:admin> sfpshow ge0
Identifier: 13 QSFP+
Connector: 12 MPO Parallel Optic
Transceiver: 000000000000004 40_Gbps \
  sw Short_dist
Encoding: 5 64B66B
Baud Rate: 103 (units 100 megabaud)
Length 9u: 0 (units km)
Length 50u (OM3): 100 (units m)
Length 62.5u:0 (units m)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-1000128-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 00000fde
Max Case Temp: 70 (C)
BR Max: 15
BR Min: 222
Serial No: LTA315281003293
Date Code: 150712
DD Type: 0x8
Enh Options: 0x0
Status/Ctrl: 0x0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 40 Centigrade
Current: 7.078 mAmps
Voltage: 3319.8 mVolts
RX Power: -2.3 dBm (593.2uW)

State transitions: 1
Last poll time: 03-17-2016 UTC Thu 08:35:18
```

```
switch:admin> sfpshow 12/ge4
Identifier: 3 SFP
Connector: 7 LC
Transceiver: 000000000000010 10_Gbps id
Encoding: 6 64B66B
Baud Rate: 103 (units 100 megabaud)
Length 9u: 0 (units km)
Length 9u: 0 (units 100 meters)
Length 50u (OM2): 8 (units 10 meters)
Length 50u (OM3): 30 (units 10 meters)
Length 62.5u:3 (units 10 meters)
Length Cu: 0 (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-0000075-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 001a Loss_of_Sig, \
```



```

Tx_Fault,Tx_Disable
BR Max:      0
BR Min:      0
Serial No:   AAF210300000PW1
Date Code:   100721
DD Type:     0x68
Enh Options: 0xf0
Status/Ctrl: 0xb0
Alarm flags[0,1] = 0x5, 0x0
Warn Flags[0,1] = 0x5, 0x0
Temperature: 42      Centigrade
Current:     8.236   mAmps
Voltage:     3333.1 mVolts
RX Power:-1.8   dBm (662.8uW)
TX Power:-2.5   dBm (561.3 uW)

State transitions: 1
Last poll time: 03-17-2016 UTC Thu 08:35:31

```

To display detailed information about a QSFP:

```

switch:admin> sfpshow 8/0
QSFP No: 0 Channel No:0
Identifier: 13 QSFP+
Connector: 12 MPO Parallel Optic
Transceiver: 0000000000000004 10_Gbps id
Encoding: 5 64B66B
Baud Rate: 103 (units 100 megabaud)
Length 9u: 0 (units km)
Length 50u (OM3): 100 (units m)
Length 62.5u:0 (units m)
Vendor Name: BROCADE
Vendor OUI: 00:17:6a
Vendor PN: 57-1000042-01
Vendor Rev: 01
Wavelength: 850 (units nm)
Options: 00000fde
Max Case Temp: 70 C)
Device Tech: 0x00
Serial No: XXXXXX16
Date Code: 100722
DD Type: 0x8
Enh Options: 0x0
Status/Ctrl: 0x0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 36      Centigrade
Current:     6.132   mAmps
Voltage:     3285.8 mVolts
RX Power:    -0.3    dBm (926.2uW)

State transitions: 1
Last poll time: 05-29-2013 UTC Wed 11:11:47

```

To display detailed information of SFP on G620:

```
switch:admin> sfpshow 0
Identifier: 3 SFP
Connector: 7 LC
Transceiver: 700c406000000000 4,8,16_Gbps \
M5,M6 sw Inter,Short_dist
Encoding: 6 64B66B
Baud Rate: 140 (units 100 megabaud)
Length 9u: 0 (units km)
Length 9u: 0 (units 100 meters)
Length 50u (OM2): 4 (units 10 meters)
Length 50u (OM3): 10 (units 10 meters)
Length 62.5u:2 (units 10 meters)
Length Cu: 0 (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-0000088-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 003a Loss_of_Sig, \
Tx_Fault,Tx_Disable
BR Max: 0
BR Min: 0
Serial No: HAA2152010C35L2
Date Code: 150517
DD Type: 0x68
Enh Options: 0xfa
Status/Ctrl: 0x80
Pwr On Time: 0.17 years (1473 hours)
E-Wrap Control: 0
O-Wrap Control: 0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 40 Centigrade
Current: 7.216 mAmps
Voltage: 3309.8 mVolts
RX Power:-2.3 dBm (583.9uW)
TX Power:-3.7 dBm (428.9 uW)

State transitions: 1
Last poll time: 03-17-2016 UTC Thu 02:08:11
```

To display information about a QSFP on a Brocade G620 switch:

```
switch:admin> sfpshow 60
QSFP No: 1 Channel No:0
Identifier: 17 QSFP28
Connector: 12 MPO Parallel Optic
Transceiver: 0802404000000080 32_Gbps OM3 sw Short_dist
Encoding: 5 64B66B
Baud Rate: 255 (units 100 megabaud)
Length 9u: 0 (units km)
Length 50u (OM4): 100 (units m)
Length 62.5u:0 (units m)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-1000331-01
Vendor Rev: A
Wavelength: 850 (units nm)
```

```

Options: 0007ffde Loss_of_Sig, Loss_of_Sig_Inverted
Max Case Temp: 0 (C)
Device Tech: 0x00
Serial No: ZTA11517000001F
Date Code: 150528
DD Type: 0xc
Enh Options: 0x0
Status/Ctrl: 0x0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 50 Centigrade
Current: 7.494 mAmps
Voltage: 3315.4 mVolts
RX Power: -inf dBm (0.0 uW)
TX Power: -25.2 dBm (3.0 uW)

State transitions: 1
Last poll time: 11-05-2015 UTC Thu 04:36:18

```

To display the cached data for a port:

```

switch:admin> sfpshow 5 -link
Identifier: 3      SFP
Connector: 7      LC
Transceiver: 540c402000000000 2,4,8_Gbps M5,M6 sw Inter_dist
Encoding: 1       8B10B
Baud Rate: 85     (units 100 megabaud)
Length 9u: 0      (units km)
Length 9u: 0      (units 100 meters)
Length 50u (OM2): 5 (units 10 meters)
Length 50u (OM3): 0 (units 10 meters)
Length 62.5u:3   (units 10 meters)
Length Cu: 0      (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-1000012-01
Vendor Rev: A
Wavelength: 850 (units nm)
Options: 001a Loss_of_Sig, Tx_Fault, Tx_Disable
BR Max: 0
BR Min: 0
Serial No: UAA1084700002C0
Date Code: 081119
DD Type: 0x68
Enh Options: 0xf0
Status/Ctrl: 0xb2
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 52   Centigrade
Current: 6.418   mAmps
Voltage: 3327.2  mVolts
RX Power: -3.9   dBm (403.5uW)
TX Power: -2.6   dBm (553.2 uW)
Port Speed Capabilities 16Gbps, 8Gbps, 4Gbps, 2Gbps

PEER PORT SFP

Optical Port: Optical Port
Optical Type: Long Wave Laser LC
Temperature: 52   Centigrade

```

sfpShow

```
Current:      6.418   mAmps
Voltage:     3327.2  mVolts
RX Power:    -3.9    dBm (403.5uW)
TX Power:    -2.6    dBm (553.2 uW)
Signal Loss (Upstream):  50 uW
Signal Loss (Downstream): 48 uW
Port Speed Capabilities 16Gbps, 8Gbps, 4Gbps, 2Gbps
```

To display the information of a remote switch port:

```
switch:admin> sfpshow -pid 0x010500
SWITCH PORT SFP

Optical Port: Optical Port
Optical Type: Long Wave Laser LC
Temperature: 52      Centigrade
Current:      6.418   mAmps
Voltage:     3327.2  mVolts
RX Power:    395.4 uW
TX Power:    544.4 uW
Port Speed Capabilities 16Gbps, 8Gbps, 4Gbps, 2Gbps

PEER PORT SFP

Optical Port: Optical Port
Optical Type: Long Wave Laser LC
Temperature: 52      Centigrade
Current:      6.418   mAmps
Voltage:     3327.2  mVolts
RX Power:    403.5 uW
TX Power:    533.2 uW
Signal Loss (Upstream):  50 uW
Signal Loss (Downstream): 48 uW
Port Speed Capabilities 16Gbps, 8Gbps, 4Gbps, 2Gbps
```

SEE ALSO [switchShow](#)

shellFlowControlDisable

Disables XON/XOFF flow control on the console serial port.

Synopsis `shellflowcontroldisable`

Description Use this command to disable XON/XOFF flow control on the console serial port. Flow control is disabled by default.

Because this command changes the flow control on the console serial port, it must be executed from a session that is logged in from the console serial port. This command cannot run from a Telnet session.

This setting is saved in the configuration database; therefore, it is persistent across reboots and power cycles.

Notes On dual-CP systems, a reboot on the standby CP is required for this command to take effect. No action is required on the active CP.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable flow control:

```
switch:admin> shellflowcontroldisable
Disabling flowcontrol
flow control is now disabled
```

SEE ALSO [shellFlowControlEnable](#)

shellFlowControlEnable

Enables XON/XOFF flow control on the console serial port.

Synopsis `shellflowcontrolenable`

Description Use this command to enable XON/XOFF flow control to the shell task. Flow control is disabled by default.

Because this command changes the flow control on the console serial port, it must be executed from a session that is logged in from the console serial port. This command cannot run from a Telnet session.

This setting is saved in the configuration database; therefore, it is persistent across reboots and power cycles.

Notes On dual-CP systems, a reboot on the standby CP is required for this command to take effect. No action is required on the active CP.

If flow control is enabled and if the console output is suspended for an extended period of time, the switch might reboot. It is recommended to disable the flow control, using **shellFlowControlDisable**.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable flow control:

```
switch:admin> shellflowcontrolenable
Enabling flowcontrol
flow control is now enabled
```

SEE ALSO [shellFlowControlDisable](#)

slotPowerOff

Removes power from a slot.

Synopsis `slotpoweroff slot [-force]`

`slotpoweroff --help`

Description Use this command to turn off the power to a blade unit. The slot must have a valid blade unit present and the blade unit must be of a type that can be powered off.

This command is supported on 8 Gbps-capable, 16 Gbps-capable, and 32 Gbps-capable core blades also. There is no frame loss when there are no ICLs connected.

This command, when executed on a core blade, displays a warning message and prompts for confirmation as powering off the last core blade will disable the chassis.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

`slot [-force]`

Specify the slot number of the blade to be powered down. This operand is required.

On core blades, the command prompts for confirmation unless you use the **-force** option.

`--help`

Displays command usage.

Examples To power off blade unit 3:

```
switch:admin> slotpoweroff 3
Slot 3 is being powered off
```

SEE ALSO [powerOffListSet](#), [powerOffListShow](#), [slotPowerOn](#), [slotShow](#)

slotPowerOn

Restores power to a slot.

Synopsis `slotpoweron slot`

`slotpoweron --help`

Description Use this command to turn on the power to a blade unit. The slot must have a valid blade unit present and the blade unit must be currently powered off. The **slotShow** command reports such slots as being in the state of INSERTED, NOT POWERED ON.

This command is supported on 8 Gbps-capable, 16 Gbps-capable, and 32 Gbps-capable core blades also. There is no frame loss when there are no ICLs connected.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

slot Specify the slot number of the blade to be powered on. This operand is required.

--help Displays command usage.

Examples To power on blade unit 3:

```
switch:admin> slotpoweron 3
Powering on slot 3.
```

SEE ALSO [slotPowerOff](#), [slotShow](#)

slotShow

Displays the status of all slots in the system.

Synopsis `slotshow [-m] [-p]`

Description Use this command to display the current status of each slot in the system. Depending on the option used, the command retrieves information on blade type, blade ID, status, Brocade model name, and power usage for each slot in the switch or chassis.

This command does not display the model names for unsupported blades. If a blade is not supported, only the Blade ID is displayed, and the Status field registers as FAULTY(9). Use the list below to identify the model name associated with an unsupported blade.

When no operand is specified, **slotShow** displays the blade type, blade ID, and status for each slot. In this view, the fields and their possible values are as follows:

| | |
|-------------------|--|
| Slot | Displays the physical slot number. |
| Blade Type | Displays the blade type as one of the following: |
| SW BLADE | The blade is a switch. |
| CP BLADE | The blade is a control processor. |
| CORE BLADE | The blade is a core switch blade. |
| AP BLADE | The blade is an application processor. |
| UNKNOWN | The blade not present or its type is not recognized. |
| ID | Displays the blade type ID as one of the following: |
| 50 | CP8 control processor blade |
| 51 | FC8-48 switch blade |
| 52 | Core8 core blade |
| 55 | FC8-32 switch blade |
| 75 | FX8-24 application processor blade |

slotShow

| | |
|-----|---------------------------------|
| 96 | FC16-48 switch blade |
| 97 | FC16-32 switch blade |
| 98 | CR16-8 core blade |
| 99 | CR16-4 core blade |
| 153 | FC16-64 switch blade |
| 175 | CPX6 control processor blade |
| 177 | CR32-8 core blade |
| 178 | FC32-48 switch blade |
| 186 | SX6 application processor blade |

Status

Displays the status of the blade as one of the following:

VACANT

The slot is empty.

INSERTED, NOT POWERED ON

The blade is present in the slot but is turned off.

POWERING UP

The blade is present and powering on.

LOADING

The blade is present, powered on, and loading the initial configuration.

DIAG RUNNING POST1

The blade is present, powered on, and running the POST (power-on self-test).

DIAG RUNNING POST2

The blade is present, powered on, and running the reboot power on self tests.

INITIALIZING

The blade is present, powered on, and initializing hardware components.

ENABLED

The blade is on and fully enabled.

DISABLED

The blade is powered on but disabled.

FAULTY

The blade is faulty because an error was detected. A fault code of 53 may also indicate the possibility of a rolling reboot detection (RRD) on the specified blade. In the case of an RRD, investigate and correct the cause, then reboot the control processor (CP) to recover the blade.

UNKNOWN

The blade is inserted but its state cannot be determined.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

In a X6-4 chassis, slot 1 and 2 (Blade ID 175) are occupied by control processor blades and slot 5 and 6 (Blade ID 176) are occupied by core blades. In a X6-8 chassis, slot 1 and 2 (Blade ID 175) are occupied by the control processor blades and slot 7 and 8 (Blade ID 177) are occupied by core blades..

Operands

This command supports the following operands:

-p

In addition to the basic slot status view, displays the following information about power consumption:

- Total direct current (DC) power consumption for the chassis and individual values for each blade (in Watts). Usage for other components such the WWN card or fans is included in the calculation but not listed per component. Note that the data displayed in the "DC Power Consumption" column and in the summary field "Total DC Power consumption" indicate the maximum allowed power consumption, not a real-time value. Real-time power consumption data is displayed for 16G platforms only with the **chassisShow** command.
- Total alternating current (AC) power consumption in Watts. This value indicates the maximum allowed AC power consumption; it is not a real-time value.
- AC efficiency, as a percentage of total and BTU.
- Power efficiency in Watts/port and Watts/Gb.

-m

In addition to the basic slot status view, displays the Brocade model name for each blade.

Examples

To display the status of all slots on a Brocade X6-8:

```
switch:user> slotshow
```

| Slot | Blade Type | ID | Status |
|------|------------|-----|---------|
| 1 | CP BLADE | 175 | ENABLED |
| 2 | CP BLADE | 175 | ENABLED |
| 3 | AP BLADE | 186 | ENABLED |
| 4 | AP BLADE | 186 | ENABLED |
| 5 | SW BLADE | 178 | ENABLED |
| 6 | SW BLADE | 178 | ENABLED |
| 7 | CORE BLADE | 177 | ENABLED |
| 8 | CORE BLADE | 177 | ENABLED |
| 9 | AP BLADE | 186 | ENABLED |
| 10 | SW BLADE | 178 | ENABLED |
| 11 | SW BLADE | 178 | ENABLED |
| 12 | AP BLADE | 186 | ENABLED |

To display the Brocade model name for each blade on a Brocade X6-8:

```
switch:user> slotshow -m
```

| Slot | Blade | Type | ID | Status |
|------|------------|------|---------|---------|
| 1 | CP BLADE | 175 | CPX6 | ENABLED |
| 2 | CP BLADE | 175 | CPX6 | ENABLED |
| 3 | AP BLADE | 186 | SX6 | ENABLED |
| 4 | AP BLADE | 186 | SX6 | ENABLED |
| 5 | SW BLADE | 178 | FC32-48 | ENABLED |
| 6 | SW BLADE | 178 | FC32-48 | ENABLED |
| 7 | CORE BLADE | 177 | CR32-8 | ENABLED |
| 8 | CORE BLADE | 177 | CR32-8 | ENABLED |
| 9 | AP BLADE | 186 | SX6 | ENABLED |
| 10 | SW BLADE | 178 | FC32-48 | ENABLED |
| 11 | SW BLADE | 178 | FC32-48 | ENABLED |
| 12 | AP BLADE | 186 | SX6 | ENABLED |

To display power consumption information on a Brocade X6-8:

```
switch:user> slotshow -p
```

| Slot | Blade Type | ID | DC Power Consumption | Status |
|------|------------|-----|----------------------|--------------------------|
| 1 | CP BLADE | 175 | 55 | ENABLED |
| 2 | CP BLADE | 175 | 55 | ENABLED |
| 3 | UNKNOWN | | - | VACANT |
| 4 | AP BLADE | 186 | 425 | INSERTED, NOT POWERED ON |
| 5 | UNKNOWN | | - | VACANT |
| 6 | SW BLADE | 178 | 260 | ENABLED |
| 7 | CORE BLADE | 177 | 450 | ENABLED |
| 8 | CORE BLADE | 177 | 450 | ENABLED |
| 9 | UNKNOWN | | - | VACANT |
| 10 | UNKNOWN | | - | VACANT |
| 11 | SW BLADE | 178 | 260 | INSERTED, NOT POWERED ON |
| 12 | UNKNOWN | | - | VACANT |

Total DC Power Consumption:
2184 watts

Total AC Power Consumption:
2426 watts AC @ 90% efficiency (8280 BTU)

Power Efficiency:
3.73 watts per port, 0.58 watts per Gb

INSERTED, NOT POWERED ON blades are not included in Total Power Consumption.

To display the Brocade model name for each blade on a Brocade X6-4:

```
switch:user> slotshow -m
```

| Slot | Blade Type | ID | Model Name | Status |
|------|------------|-----|------------|---------|
| 1 | CP BLADE | 175 | CPX6 | ENABLED |
| 2 | CP BLADE | 175 | CPX6 | ENABLED |
| 3 | SW BLADE | 178 | FC32-48 | ENABLED |
| 4 | SW BLADE | 178 | FC32-48 | ENABLED |

| | | | | |
|---|------------|-----|---------|---------|
| 5 | CORE BLADE | 176 | CR32-4 | ENABLED |
| 6 | CORE BLADE | 176 | CR32-4 | ENABLED |
| 7 | SW BLADE | 178 | FC32-48 | ENABLED |
| 8 | AP BLADE | 186 | SX6 | ENABLED |

To display power consumption information on a Brocade X6-4:

```
switch:user> slotshow -p
Slot   Blade Type      ID    DC Power      Status
        Consumption
-----
  1     CP BLADE      175     55     ENABLED
  2     UNKNOWN
  3     UNKNOWN
  4     UNKNOWN
  5     CORE BLADE   176    250     ENABLED
  6     UNKNOWN
  7     SW BLADE     178    260     ENABLED
  8     SW BLADE     178    260     ENABLED
Total DC Power Consumption:
1435 watts
Total AC Power Consumption:
1594 watts AC @ 90% efficiency (5440 BTU)
Power Efficiency:
4.15 watts per port, 0.26 watts per Gb
```

SEE ALSO [bladeDisable](#), [bladeEnable](#), [chassisShow](#), [slotPowerOff](#), [slotPowerOn](#)

snmpConfig

Manages the SNMP agent configuration.

Synopsis

```
snmpconfig --set snmpv1 [-index index] [-community community_name]
  [-host [ipv4 | ipv6 | dns]] [-port udp_port]
  [-severity severity_level]

snmpconfig --set snmpv1 [-index index] [-default]

snmpconfig --set snmpv3 [-enable | -disable [informs]]

snmpconfig --set snmpv3 [-index index] [-user user_name]
  [-host [ipv4 | ipv6 | dns]] [-port udp_port]
  [-severity severity_level] [-auth_proto auth_protocol]
  [-auth_passwd auth_password] [-priv_proto priv_protocol]
  [-priv_passwd priv_password] [-engine_id engine_id]
  [-notify_type trap | informs]

snmpconfig --set mibcapability [-mib_name mib_name [-bitmask bit_mask]]

snmpconfig --set accesscontrol [-index index] [-host ip_address]
  [-access [ro | rw]]

snmpconfig --set systemgroup [-sysdescr sysdescr] [-syslocation syslocation]
  [-syscontact syscontact] [-authtrapeenabled [ON | OFF]]

snmpconfig --set seclevel [-snmpget security_level] [-snmpset security_level]

snmpconfig --set auditinterval [-interval minute]

snmpconfig --enable | --disable mibcapability
  -mib_name mib_name [-trap_name trap_name]

snmpconfig --enable | --disable snmpv1

snmpconfig --show [snmpv1 | snmpv3 | accesscontrol | mibcapability |
  systemgroup | seclevel | auditinterval]

snmpconfig --help
```

Description

Use this command to manage the configuration of the SNMP agent in the switch. The configuration includes SNMPv1 and SNMPv3 configuration, access control list (ACL), MIB capability, system group, and security level settings. The command supports set, reset to default, and display operations. This command supports enabling or disabling a single MIB or all MIBs, configuring a single trap only, and managing traps in excess of 32.

The SNMP Agent configuration interface allows both interactive and command line options (non-interactive) for all parameters.

The SNMPv3 configuration supports sending inform requests as an alternative to trap requests. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the manager does not receive an inform request, it does not send a response. If the sender never receives a response, the inform request can be sent again. For this reason, informs are more likely to reach their intended destination.

All values successfully changed by this command take effect immediately and are persistent across power cycles and reboots.

SNMPv1 Configuration Parameters

The agent supports six communities and their associated trap recipients and trap recipient severity levels. The first three communities are for read-write (rw) access and the last three are for read-only (ro) access. The default value for the trap recipient of each community is 0.0.0.0. The length of the community string should be in the range of 2 to 16 characters. The default values for the community strings include the following:

- Community 1: Secret Code
- Community 2: OrigEquipMfr
- Community 3: private
- Community 4: public
- Community 5: common
- Community 6: FibreChannel

For an SNMP management station to receive a trap generated by the agent, the administrator must configure a trap recipient to correspond to the IP address of the management station. In addition, the trap recipient must be able to pass the access control list (ACL) check as described in the Access Control category.

SNMPv3 Configuration Parameters

Two user roles, **snmpadmin** and **snmpuser** are supported. The **snmpadmin** role provides read-write access and the **snmpuser** role provides read-only access. Entries are added to the USM table corresponding to each role. A total of three entries for **snmpadmin** and **snmpuser** respectively are supported. Separate default passwords are provided for creation of **authKey** and **privKey** for each entry. The default set of passwords is published and the default algorithm (MD5/SHA) is used to create the initial set of authentication keys. You can change these passwords using this option. You can select the authentication protocol MD5/SHA or no authentication for each entry.

When new passwords are entered for any user entry, a new **authKey** and **privKey** are generated. The new passwords must be updated on the client (e.g., MIB browser) as well. **AuthKey** and **privKey** can also be updated with the delta key mechanism provided by the SNMPv3 protocol.

The system prompts for password confirmation if a protocol other than **NoAuth/NoPriv** is selected. Protocol passwords must be between 1 and 20 characters.

In order for an SNMP management station to receive SNMPv3 traps generated by the agent, the administrator must configure a trap recipient value to correspond to the IP address of the management station. In addition, the trap recipient must pass the ACL check as described in the Access Control section. The trap recipient must be associated with one of the six users of SNMPv3 and trap severity level. The factory default value for the SNMPv3 trap recipient of each user is 0.0.0.0.

Access Control Configuration Parameters

The ACL check is as follows: there are six ACLs to restrict SNMP get, set, and trap operations to hosts under an host-subnet-area. The host-subnet-area is defined by comparing nonzero IP octets. For example, an ACL of 192.168.64.0 enables access by any hosts that start with the specified octets. The connecting host is enabled to set each host-subnet-area to be read-write or read-only. The closest match out of six entries is given access. The ACL check is turned off when all six entries contain 0.0.0.0. The default values of all six entries are 0.0.0.0. For IPv6 subnets, the format is specified by an IPv6 address followed by the number of fixed bits in the address.

MIB Capability Configuration Parameters

The **mibCapability** option turns certain MIBs and associated SNMP traps on or off. If a specific MIB is disabled, the corresponding traps also are disabled. If any trap group is disabled, the corresponding individual traps are also disabled.

SNMP Traps are identified by their bit mask and can be read directly from the switch configuration. The MIB and trap status (enabled or disabled) status is recorded in a 64-bit counter. The last bit (bit 0) is reserved for the MIB and the remaining bits are reserved for the traps of that MIB. The trap's position is allocated based on the last ID of the trap OID. For example, the last ID of the swEventTrap is 5 so its position will be 5th from the right. The following is a listing of valid SNMP traps and their bit masks:

| MIB | Trap Name and position | Bit mask | Default |
|-----------|---------------------------------|----------|----------|
| FE-MIB | | 0x1 | Enabled |
| SW-MIB | | 0x1 | Enabled |
| | swFault (1) | 0x2 | Enabled |
| | swSensorScn (2) | 0x4 | Enabled |
| | swFCPortScn (3) | 0x8 | Enabled |
| | swEventTrap (4) | 0x10 | Enabled |
| | swIPv6ChangeTrap (7) | 0x80 | Enabled |
| | swPmgrEventTrap (8) | 0x100 | Enabled |
| | swFabricSegmentTrap (9) | 0x200 | Disabled |
| | swFabricReconfigTrap (10) | 0x400 | Disabled |
| | swExtTrap (11) | 0x800 | Disabled |
| | swStateChange (12) | | Disabled |
| | swMovePort (13) | | Disabled |
| | swBrcdGenericTrap (15) | | Disabled |
| | swDeviceStatusTrap (16) | | Disabled |
| | swZoneConfigChangeTrap (17) | | Disabled |
| FA-MIB | | 0x1 | Enabled |
| | connUnitStatusChange(1) | 0x2 | Enabled |
| | connUnitAddedTrap (2) | N/A | Enabled |
| | connUnitDeletedTrap (3) | 0x8 | Enabled |
| | connUnitEventTrap (4) | 0x10 | Enabled |
| | connUnitSensorStatusChange(5) | 0x20 | Enabled |
| | connUnitPortStatusChange (6) | 0x40 | Enabled |
| FICON-MIB | | 0x1 | Enabled |
| | linkRNIDDeviceRegistration (1) | 0x2 | Enabled |
| | linkRNIDDevicedeRegistration(2) | 0x4 | Enabled |
| | linkLIRRListenerAdded (3) | 0x8 | Enabled |
| | linkLIRRListenerRemoved (4) | 0x10 | Enabled |
| | linkRLIRFailureIncident (5) | 0x20 | Enabled |
| HA-MIB | | 0x1 | Enabled |

| | | |
|------------------------------|------|----------|
| fruStatusChanged (1) | 0x2 | Enabled |
| cpStatusChanged (2) | 0x4 | Enabled |
| fruHistoryTrap (3) | 0x8 | Enabled |
| <hr/> | | |
| FCIP-MIB | 0x1 | Enabled |
| <hr/> | | |
| IF-MIB | 0x1 | Enabled |
| linkUpTrap (3) | 0x8 | Enabled |
| linkdownTrap (4) | 0x10 | Enabled |
| <hr/> | | |
| MAPS-MIB | 0x1 | Enabled |
| mapsTrapAM (1) | 0x2 | Enabled |
| <hr/> | | |
| T11-FC-ZONE-SERVER-MIB | 0x1 | Disabled |
| t11ZsRequestRejectNotify (1) | 0x2 | Disabled |
| t11ZsMergeFailureNotify (2) | 0x4 | Disabled |
| t11ZsMergeSuccessNotify (3) | 0x8 | Disabled |
| t11ZsDefZoneChangeNotify (4) | 0x10 | Disabled |
| t11ZsActivateNotify (5) | 0x20 | Disabled |

Note: The Zone Mib parameters are for restricted usage only.

Use the **--show mibcapability** option to display the traps configurable under each MIB.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Informs are not supported for IPv6 addresses.

Operands

This command has the following operands:

--help

Displays the command usage.

--enable

Enables the SNMP agent configuration for the specified category. This operand is valid only with **mibcapability** and **snmpv1**. When used with the **snmpv1** operand, this command restores access to SNMPv1/v2c.

--disable

Disables the SNMP agent configuration for the specified category. This operand is valid only with **mibcapability** and **snmpv1**. When used with the **snmpv1** operand, this command blocks access to SNMPv1/v2c. All requests for v1/v2c version will be dropped by the switch, and SNMPv1 traps will be blocked from being sent, even if trap destinations are present.

--show

Displays the SNMP agent configuration data of the specified category. When used with the **snmpv1** operand, this command displays whether access to SNMPv1/v2c is enabled or disabled.

--set

Sets the SNMP agent configuration data of the specified category.

snmpv1

Selects SNMPv1-related configuration parameters. These parameters include the community string, trap recipient IP address, and trap severity level associated with each trap recipient IP address. When "0" is configured as a trap port, traps can be received at the default port 162.

snmpv3

Selects SNMPv3-related configuration parameters. These parameters include the user name, authentication protocol and password, the privacy protocol and password, the SNMPv3 trap recipient's IP address, its associated user index, and trap severity level. When "0" is configured as a trap port, traps can be received at the default port 162.

-index *index*

Specifies the index for which the user details are configured. Valid values are from 1 through 6.

-community *community_name*

Specifies the community string. If special characters are used in the *community_name*, the *community_name* must be enclosed in single quotes. The default values for the community strings include the following:

- Community 1: Secret Code
- Community 2: OrigEquipMfr
- Community 3: private
- Community 4: public
- Community 5: common
- Community 6: FibreChannel

-host [*ipv4* | *ipv6* | *dns*]

Specifies the IP address of the host. IPv4, IPv6, and DNS hosts are supported.

-port *udp_port*

Specifies the UDP port where SNMP traps will be received. Valid port IDs range from 0 through 65535. The default port is 162.

-severity *severity_level*

Specifies the trap recipient severity level. When an event occurs and its severity level is at or below the set value, the Event Trap traps (swEventTrap and connUnitEventTrap), are sent to configured trap recipients. By default, this value is set at 0, implying that no Event Trap is sent. Possible values are:

- 0: None
- 1: Critical
- 2: Error
- 3: Warning
- 4: Informational
- 5: Debug

--default

Sets the SNMP agent configuration data for a specified item to the default values. Generally, these default values may be available in the configuration database. The command sets to factory defaults if the SNMP agent configuration parameters are not available in the configuration database. This operand is valid only with **snmpv1**.

-enable | -disable [informs]

Enables or disables informs. If informs are enabled, all trap destinations receive inform requests. If informs are disabled, all trap destinations receive trap requests. When informs are enabled, the engine ID must be set to correspond to the management engine IP address. Informs are by default disabled. IPv6 Informs are currently not supported.

-user *user_name*

Specifies the name of the user that connects to the agent. The user name must be between 2 and 32 characters long. The default user names are defined with the **noAuth** and **noPriv** protocol. The factory default SNMPv3 user names include the following:

- User 1: snmpadmin1
- User 2: snmpadmin2
- User 3: snmpadmin3
- User 4: snmpuser1
- User 5: snmpuser2
- User 6: snmpuser3

-auth_proto *auth_protocol*

Specifies the authentication protocol as MD5, SHA, or no authentication for each entry. The following combinations of protocols are supported:

- NoAuth/NoPriv
- Auth/NoPriv
- Auth/Priv

-auth_passwd *auth_password*

Specifies the authentication password that enables the agent to receive packets from the host. Passwords are plain text and must be added each time for each configuration replay. The password must be between 1 and 32 characters long.

-priv_proto *priv_protocol*

Specifies privacy authentication level setting. The valid values are **DES** and **AES128**.

-priv_passwd *priv_password*

Specifies the privacy password that enables the host to encrypt the contents of the message that it sends to the agent. Passwords are plain text and must be added each time for each configuration replay. The privacy password alone cannot be configured. You configure the privacy password with the authentication password.

-engine_id *engine_id*

Configures a user-defined engine ID for the SNMP agent.

-notify_type trap | informs

Specifies the type of notification traps that are sent for the host. Traps and informs are supported. The default notify type is traps.

--set mibcapability

Configures MIBs.

-mib_name *mib_name*

Specifies the name of the MIB to be configured. This operand is required if you want to configure MIB traps noninteractively. Valid MIB names include the following.

- FE-MIB
- SW-MIB
- FA-MIB
- FICON-MIB
- HA-MIB
- FCIP-MIB
- ISCSI-MIB
- IF-MIB
- MAPS-MIB
- T11-FC-ZONE-SERVER-MIB

-bitmask *bit_mask*

Specifies the bit mask for the MIB. Refer to the table above for valid values.

-trap_name *trap_name*

Specifies the name of the trap to be enabled or disabled. This operand is optional. Use **snmpconfig --show mibcapability** for a listing of valid traps.

accessControl

Selects access-control-related parameters. These parameters include the access host subnet area and access permission (read-write).

-host *ip_address*

Specifies the host IP address with or without a subnet mask.

-access [ro | rw]

Specifies the read-write (rw) or read-only (ro) access to the host.

systemGroup

Selects configuration parameters related to the system group.

-sysdescr *sysdescr*

Specifies the system description. The default value is "Fibre Channel Switch".

-syslocation *syslocation*

Specifies the location of the system (switch). The default value is "End User Premise". The string must be at least 4 characters in length; the maximum length is 255 characters. If special characters are used in the *syslocation*, the *syslocation* must be enclosed in single quotes.

-syscontact *syscontact*

Specifies the contact information for this system (switch). The default value is "Field Support". Refer to the definition of sysDescr, sysLocation and sysContact in the system group of MIB-II. The string must be at least 4 characters in length; the maximum length is 255 characters. If special characters are used in the *syscontact*, the *syscontact* must be enclosed in single quotes.

-authtrapenabled [ON | OFF]

Enables or disables the authentication traps. When enabled, the authentication trap (authenticationFailure) is transmitted to a configured trap recipient in the event that the agent receives a protocol message that is not properly authenticated. In the context of SNMPv1 and SNMPv2c, this means that a request contains a community string that is not known to the agent. The default value for this parameter is OFF (disabled).

secLevel

Sets the SNMP security level.

-snmpget *security_level*

Specifies security level for all SNMP GET requests.

-snmpset *security_level*

Specifies security level for SNMP SET requests only. Values include the following:

- 0 - No security.
- 1 - Authentication only.
- 2 - Authentication and Privacy.
- 3 - OFF

auditinterval

Sets the SNMP audit interval.

-interval *minute*

Specifies the the SNMP audit interval in minutes. Valid value range is from 1 through 1440. Default value is 60.

Examples

To display the SNMPv1 configuration :

```
switch:admin> snmpconfig --show snmpv1
SNMPv1 community and trap recipient configuration:
Community 1: Secret C0de (rw)
  Trap recipient: 10.32.147.113
  Trap recipient Severity Level: 0
Community 2: OrigEquipMfr (rw)
  Trap recipient: 1080::8:800:200C:1234
  Trap recipient Severity Level: 0
Community 3: private (rw)
  No trap recipient configured yet
Community 4: public (ro)
  No trap recipient configured yet
Community 5: common (ro)
  No trap recipient configured yet
Community 6: FibreChannel (ro)
  No trap recipient configured yet
```

To set the SNMPv1 configuration of a switch noninteractively:

```
switch:admin> snmpconfig --set snmpv1 -index 2 -community community2 \
  -port 4000 -severity 3
Committing configuration.....done
```

To set the SNMPv3 configuration of a switch noninteractively:

```
switch:admin> snmpconfig --set snmpv3 -index 2 -user user1 \
    -host 1.3.4.5 -port 178 -severity 5 -auth_proto 2 \
    -auth_passwd "password1" -priv_proto 3 -priv_passwd "password2" \
    -engine_id "80:00:08:08:08:98:AB:5C:01" -notify_type informs \
    Committing configuration.....done
```

To set the severity level for switch events and MAPS alerts interactively:

```
switch:admin> snmpconfig --set mibcapability
[...]
SW-TRAP (yes, y, no, n): [yes]
swFCPortScn (yes, y, no, n): [yes]
swEventTrap (yes, y, no, n): [yes]

    DesiredSeverity: (0..4) [0] 4
swIPv6ChangeTrap (yes, y, no, n): [yes]
    DesiredSeverity: (0..4) [0] 4
swPmgrEventTrap (yes, y, no, n): [yes]
[...]
```

To enable the **mapsTrapAM** noninteractively:

```
switch:admin> snmpconfig --enable mibcapability \
    -mib_name SW-MIB -trap_name mapsTrapAM
Operation succeeded
```

To enable the **swEventTrap** of the SW-MIB category only (this operation disables all other SNMP traps in this MIB category):

```
switch:admin> snmpconfig --set mibcapability \
    -mib_name SW-MIB -bitmask 0x10
Operation succeeded

switch:admin> snmpconfig --show mibcapability
[...]
SW-TRAP: NO
    swFCPortScn: YES
    swEventTrap: YES
        DesiredSeverity:4
    swIPv6ChangeTrap: YES
        DesiredSeverity:None
    swPmgrEventTrap: YES
    swFabricReconfigTrap: YES
    swFabricSegmentTrap: YES
    swExtTrap: NO
    swStateChangeTrap: NO
    swPortMoveTrap: NO
    swBrcdGenericTrap: YES
    swDeviceStatusTrap: YES
    swZoneConfigChangeTrap: NO
[...]
```

To enable the SW-MIB MIB only without changing the current trap configuration:

```
switch:admin> snmpconfig --enable mibcapability \
               -mib_name SW-MIB
Operation succeeded

switch:admin> snmpconfig --show mibcapability
[...]
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
        DesiredSeverity:4
    swIPv6ChangeTrap: YES
        DesiredSeverity:4
    swPmgrEventTrap: YES
    swFabricReconfigTrap: YES
    swFabricSegmentTrap: YES
    swExtTrap: NO
    swStateChangeTrap: NO
    swPortMoveTrap: NO
    swBrcdGenericTrap: YES
    swDeviceStatusTrap: YES
    swZoneConfigChangeTrap: NO
[...]
```

To re-enable all traps under the SW-MIB category:

```
switch:admin> snmpconfig --set mibcapability \
               -mib_name SW-MIB -bitmask OxFFF
Operation succeeded

switch:admin> snmpconfig --show mibcapability
[...]
SW-TRAP: YES
    swFault: YES
    swSensorScn: YES
    swFCPortScn: YES
    swEventTrap: YES
        DesiredSeverity:None
    swIPv6ChangeTrap: YES
    swPmgrEventTrap: YES
    swFabricReconfigTrap: Yes
    swFabricSegmentTrap: Yes
    swExtTrap: Yes
    swStateChangeTrap: NO
    swPortMoveTrap: NO
    swBrcdGenericTrap: NO
[...]
```

To display the configuration for all MIBs and associated traps:

```
switch:admin> snmpconfig --show mibcapability
FE-MIB: YES
SW-MIB: YES
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
FCIP-MIB: YES
IF-MIB: YES
```

```

BD-MIB: YES
BROCADE-MAPS-MIB: YES
T11-FC-ZONE-SERVER-MIB: NO
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
        DesiredSeverity:None
    swIPv6ChangeTrap: YES
        DesiredSeverity:None
    swPmgrEventTrap: YES
    swFabricReconfigTrap: YES
    swFabricSegmentTrap: YES
    swExtTrap: NO
    swStateChangeTrap: NO
    swPortMoveTrap: NO
    swBrcdGenericTrap: YES
    swDeviceStatusTrap: YES
    swZoneConfigChangeTrap: NO
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitEventTrap: YES
    connUnitPortStatusChange: YES
FICON-TRAP: YES
    linkRNIDDeviceRegistration: YES
    linkRNIDDeviceDeRegistration: YES
    linkLIRRListenerAdded: YES
    linkLIRRListenerRemoved: YES
    linkRLIRFailureIncident: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
IF-TRAP: YES
    linkDown: YES
    linkUp: YES
BD-TRAP: YES
    bdTrap: YES
    bdClearTrap: YES
MAPS-TRAP: YES
    mapsTrapAM: YES
T11-FC-ZONE-SERVER-TRAP: NO
    t11ZsRequestRejectNotify: NO
    t11ZsMergeFailureNotify: NO
    t11ZsMergeSuccessNotify: NO
    t11ZsDefZoneChangeNotify: NO
    t11ZsActivateNotify: NO

```

Note: The Zone Mib parameters are for restricted usage only.

To set the access control configuration noninteractively:

```

switch:admin> snmpconfig --set accesscontrol -index 1 -host 2.3.4.5 -access ro
Committing configuration...done.

```


To display the access control configuration:

```
switch:admin> snmpconfig --show accesscontrol
SNMP access list configuration:
Entry 0: Access host subnet area 2.3.4.5 (ro)
Entry 1: No access host configured yet
Entry 2: No access host configured yet
Entry 3: No access host configured yet
Entry 4: No access host configured yet
Entry 5: No access host configured yet
```

To configure system group parameters and verify the configuration noninteractively:

```
switch:admin> snmpconfig --set systemgroup -syscontact "Field Support" \
  -authTrapEnabled false -sysdescr "Fibre Channel Switch" \
  -syslocation "End User Premise"

switch:admin> snmpconfig --show systemgroup
sysDescr = Fibre Channel Switch
sysLocation = End User Premise
sysContact = Field Support
authTraps = 0 (OFF)
```

To set the SNMP security level and verify the configuration noninteractively:

```
switch:admin> snmpconfig --set seclevel -snmpset 0
SET seclevel cannot be less than GET. Changing SET seclevel to be \
  same as GET. Do you want to continue? (yes, y, no, n): [no] y

switch:admin> snmpconfig --show seclevel
GET security level = 1, SET level = 1
SNMP GET Security Level: Authentication only
SNMP SET Security Level: Authentication only
```

To set the audit interval and verify the configuration noninteractively:

```
switch:admin> snmpconfig --set auditinterval -interval 90
Committing configuration....done.

switch:admin> snmpconfig --show auditinterval
SNMP Audit Interval (in min): 90
```

To display the SNMPv3 configuration with informs disabled interactively:

```
switch:admin> snmpconfig --show snmpv3

SNMP Informs = 0 (OFF)

SNMPv3 USM configuration:
User 1 (rw): snmpadmin1
    Auth Protocol: noAuth
    Priv Protocol: noPriv
User 2 (rw): snmpadmin2
    Auth Protocol: noAuth
    Priv Protocol: noPriv
User 3 (rw): snmpadmin3
    Auth Protocol: noAuth
    Priv Protocol: noPriv
User 4 (ro): snmpuser1
```

```

        Auth Protocol: noAuth
        Priv Protocol: noPriv
User 5 (ro): snmpuser2
        Auth Protocol: noAuth
        Priv Protocol: noPriv
User 6 (ro): admin
        Auth Protocol: noAuth
        Priv Protocol: noPriv

```

```

SNMPv3 Trap/Informs configuration:
Trap Entry 1:      1.2.3.4
    Trap Port: 162
    Trap User: snmpadmin1
    Trap recipient Severity level: 4
    Notify Type: TRAP(1)
Trap Entry 2:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 3:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 4:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 5:      No trap recipient configured yet
    Notify Type: TRAP(1)
Trap Entry 6:      5.6.7.8
    Trap Port: 162
    Trap User: admin
    Trap recipient Severity level: 4
    Notify Type: TRAP(1)

```

To enable inform requests to be sent instead of trap requests interactively:

```

switch:admin> snmpconfig --set snmpv3

SNMP Informs Enabled (true, t, false, f): [true]

User (rw): [snmpadmin1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [01:02:03:04:05:06:07:08:09]
User (rw): [snmpadmin2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [00:00:00:00:00:00:00:00:00]
User (rw): [snmpadmin3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [00:00:00:00:00:00:00:00:00]
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [00:00:00:00:00:00:00:00:00]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [00:00:00:00:00:00:00:00:00]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [00:00:00:00:00:00:00:00:00]

```

```

SNMPv3 trap/inform recipient configuration:
Notify Type [TRAP(1)/INFORM(2)]:(1..2)[1] 2
Inform Recipient's IP address : [0.0.0.0] 10.70.12.115
UserIndex: (1..6) [1]
Inform recipient Severity level : (0..5) [0] 5
Inform recipient Port : (0..65535) [162]
Notify Type [TRAP(1)/INFORM(2)]:(1..2)[1] 2
Inform Recipient's IP address : [0.0.0.0] 1.2.3.4
UserIndex: (1..6) [2]
Informs recipient Severity level : (0..5) [0] 2
Informs recipient Port : (0..65535) [162]
Notify Type [TRAP(1)/INFORM(2)]:(1..2)[1]
Trap Recipient's IP address : [0.0.0.0]
Notify Type [TRAP(1)/INFORM(2)]:(1..2)[1]
Trap Recipient's IP address : [0.0.0.0]
Notify Type [TRAP(1)/INFORM(2)]:(1..2)[1] 1
Trap Recipient's IP address : [0.0.0.0] 10.70.4.106
UserIndex: (1..6) [5]
Trap recipient Severity level : (0..5) [0] 2
Trap recipient Port : (0..65535) [162]
Notify Type [TRAP(1)/INFORM(2)]:(1..2)[1]
Trap Recipient's IP address : [0.0.0.0]
Committing configuration.....done.

```

To block access to SNMPv1/2c

```

switch:admin> snmpconfig --disable snmpv1
switch:admin> snmpconfig --show
SNMPv1 community and trap recipient configuration:
Community 1: ram (rw)
    Trap recipient: 172.26.4.84
    Trap port: 162
    Trap recipient Severity level: 5
Community 2: OrigEquipMfr (rw)
    No trap recipient configured yet
Community 3: private (rw)
    No trap recipient configured yet
Community 4: public (ro)
    No trap recipient configured yet
SNMPv1:Disabled

```

References Refer to the following publications for further information on SNMP:

- *SW_v5_x.mib, "Switch Management Information & Switch Enterprise Specific Trap"*
- *RFC1157, "A Simple Network Management Protocol (SNMPv1)"*
- *RFC1213, "Management information Base for Network Management of TCP/IP-based internets: MIB-II"*
- *RFC2574, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)"*

SEE ALSO None

snmpTraps

Sends or displays SNMP traps.

Synopsis `snmptraps --send [-trap_name trap_name]
[-ip_address ip_address]`

`snmptraps --show [port]`

`snmptraps --block -port [slot]port | ALL`

`snmptraps --unblock -port [slot]port | ALL`

`snmptraps --help`

Description Use this command to manage specific Simple Network Management Protocol (SNMP) traps.

Use the **--send** option to send a specific SNMP trap to a recipient indicated by its IP address. Or use the **--send** option without operands to send all supported traps to all configured SNMP trap recipients. When the command is issued to send all traps, the message returned indicates only the total number of traps sent and not the individual trap names.

In Fabric OS v7.0.0 and later you can block or unblock certain port traps on specified ports. This feature provides control over a subset of port traps including the following: SwFCPortScn, ConnUnitPortStatusChange, and SwFabricSegmentTrap.

Use the **--show** option with the **port** operand to display the status of blocked ports. When used without operands, the **--show** option displays all Management Information Base (MIB) objects and associated traps that are supported in Fabric OS.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--send

Sends one or all SNMP traps to all configured recipients or to a specified recipient. The following operands are optional:

-trap_name *trap_name*

Specifies the trap by name. Use `snmptraps --show` for a listing of valid traps.

-ip_address *ip_address*

Specifies the recipient by its IP address in IPv4 or IPv6 format. IPv6 addresses require Fabric OS v6.4.0 or later.

--block [*slot*]*port* | ALL

Blocks the following port traps on the specified port: SwFCPortScn, ConnUnitPortStatusChange, and SwFabricSegmentTrap. When used with the **ALL** this command blocks these traps on all ports.

--unblock [*slot*]*port* | ALL

Removes the trap blocking from the specified port or from all ports.

--show [port]

Displays all configured SNMP traps and MIBs. When used with the optional **port** argument, this command displays the ports that are blocked from receiving certain software traps.

--help

Displays the command usage.

Examples

To send all traps to the configured recipients:

```
switch:admin> snmpTraps --send
Number of traps sent : 27
```

To send a FRU history trap to recipient 172.16.0.12

```
switch:admin> snmptraps --send -trap_name \
    fru-history-trap -ip_address 172.16.0.12.
Number of traps sent : 1
```

To display the traps and MIBs supported in Fabric OS:

```
switch:admin> snmpTraps --show
```

| # | Mib Name | Supported Traps |
|-----|------------------------|---|
| 001 | SW-MIB | sw-fc-port-scn sw-ip-v6-change-trap sw-pmgr-event-trap sw-event-trap sw-fabric-reconfig-trap sw-fabric-segment-trap sw-state-trap sw-zone-config-change-trap sw-port-move-trap sw-brcd-generic-trap sw-device-status-trap |
| 002 | FICON-MIB | link-rnid-device-registration link-rnid-device-deregistration link-lirr-listener-added link-lirr-listener-removed link-rlir-failure-incident |
| 003 | FA-MIB | conn-unit-status-change conn-unit-sensor-status-change conn-unit-port-status-change conn-unit-event-trap |
| 004 | RFC1157 | cold-restart-trap warm-restart-trap if-link-up-trap if-link-down-trap snmp-authetication-trap |
| 005 | HA-MIB | fru-status-change-trap fru-history-trap cp-status-change-trap |
| 006 | T11-FC-ZONE-SERVER-MIB | t11ZsRequestRejectNotify t11ZsMergeSuccessNotify t11ZsMergeFailureNotify t11ZsDefZoneChangeNotify t11ZsActivateNotify |

snmpTraps

Note: The Zone Mib parameters are for restricted usage only.

To block a single port from receiving traps:

```
switch:admin> snmpTraps --block 17
```

To unblock a previously blocked port:

```
switch:admin> snmpTraps --unblock 17
```

To configure the switch so that only two ports receive the port traps, you block all traps and then unblock the ports you want to receive the traps:

```
switch:admin> snmpTraps --block ALL
```

```
switch:admin> snmpTraps --unblock 17
```

```
switch:admin> snmpTraps --unblock 18
```

To display the blocked port status:

```
switch:admin> snmpTraps --show port
Port      0      1      2      3      4      5      6      7      8      9
-----
Blocked   |      |      |      |      |      |      |      |      |
Port     10     11     12     13     14     15     16     17     18     19
-----
Blocked   |      |      |      |      |      |      | Y | Y |
Port     20     21     22     23     24     25     26     27     28     29
-----
Blocked   |      |      |      |      |      |      |      |      |
Port     30     31     32     33     34     35     36     37     38     39
-----
Blocked   |      |      |      |      |      |      |      |      |
```

SEE ALSO [snmpConfig](#)

spinFab

Runs functional test of interswitch link (ISL) cabling and trunk group operation.

Synopsis

```
spinfab
  [-nmegs count]
  [-framesize bytes]
  [-ports itemlist]
  [-setfail mode]
  [-fports flag]
  [-nframes number]
  [-pattern number/name]
  [-timeout length]
```

Description

Use this command to verify the intended functional operation of interswitch links (ISLs) at the maximum speed by setting up the routing hardware so that test frames received by each E_Port are retransmitted on the same E_Port. Several frames are subsequently sent to the neighbor port attached to each active E_Port specified. Because the default action for such frames is to route them back to the sender, which never occurs during normal traffic, the frames circulate until the test terminates.

The frames are continuously transmitted and received in all ports in parallel. The port LEDs flicker green rapidly while the test is running.

M->N/M->M loopback ports are tested as well, using the same algorithm, if loopback cables or loopback plugs are present in the switch.

While the frames are circulating, the RX frame count and port CRC and encoder error statistics are monitored. If a port stops or a low-level error occurs, the test generates an error message. Every one million frames, the circulating frames are captured to verify that they are still circulating and in the appropriate order. In this manner, the test can verify the entire path to the remote switch as well as the proper in-order delivery operation of any trunk groups present.

In case of trunk master ports, all the slave ports are also monitored for low-level errors.

The switch remains in normal operation while this test is running. However, some performance degradation may occur due to the ISLs being saturated with test frames. For this reason, use caution when running this test on live fabrics. Consider testing only one trunk group or ISL at a time, and do not run the tests for extended periods of time.

Combine this test with **portLoopBackTest** for ISL link failure isolation. If **spinFab** fails, replace the cable with a loopback plug and run **portLoopBackTest** to verify the local switch and media. If these pass, the fault lies in the cable, the remote switch, or media.

The frame size depends on the amount of buffer credit available on the port. There are eight possible frames that can be sent. Especially with trunking groups, all eight possible frames are used unless there is extensive traffic running on the link. The payload sizes of those eight frames are 1024, 12, 8,1024, 512, 1024, 12, and 1024.

Notes The following port types support the **spinFab** diagnostics.

- Loopback ports
- AE_Ports
- E_Ports
- Interchassis link (ICL) E_Ports
- ICLs configured as D_Ports
- D_Ports
- F_Ports connected to Brocade-branded HBA
- Trunk master ports
- Trunk slave ports
- Ports in logical switches
- Ports in Base Switches
- Loopback ports in D_Port mode between two logical switches on the same switch
- Ports on non-VF connected to VF-configured switch
- Long distance ports
- D_Ports connected to AG switch in D_Port mode
- ICL D_Ports connected to AG switch in D_Port mode
- Loopback ports in D_Port mode connected to AG switch

The following ports do not support the **spinFab** test:

- F_Ports connected to non-Brocade-branded HBA
- F_Port connections to N_Port on AG switch
- F_Port on VF-enabled switch
- Ports on AG switch connected to HBA
- EX_Ports
- E_Ports connected to EX_Ports
- Core blade connected to port blade

When trunk groups are present, the entire trunk group must be included in the range of ports to test or false failure notifications may occur. If multiple ISL links are present between two switches that support trunking, then it is likely that trunk groups are present and all ports between the two switches should be tested at the same time.

You cannot interrupt the test by pressing the return key (<cr>).

This command supports a maximum of 16 paths to a single remote domain.

When new logical switches are created, appropriate switch domain must be set to avoid domain ID overlap when running spinfab. With insistent DID turned off to overcome this issue, DID negotiation depends on when the switch is powered on. A reboot of the switches with links already established re-configures the domain if the domain ID is already in use by another switch in the fabric, to avoid convergence.

Spinfab tests skips F_Ports if Virtual Fabric is enabled on any platform.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

-nmegs *count*

Specifies the number of frames to send in millions. The test progresses until the specified number of frames has been transmitted on each port. The default value is 10 million frames. This command only approximately counts the frames and the actual number of frames sent will be slightly larger, particularly at link speeds of 4 Gbps or higher.

-framesize *bytes*

For internal use only. Not supported.

-ports *itemlist*

Specifies a list of user ports to test. By default, all of the ISL ports in the current switch are tested. Refer to **itemList** for further details.

-setfail *mode*

Instructs **spinFab** how to mark failed ports. Valid values include the following:

0

Does not mark failing ports as FAILED (default). This option minimizes the impact on live fabrics.

1

Marks the failing ports as FAILED. In test or qualification environments without live traffic, this may be useful with large values of **-nmegs count**. This mode is disabled by default.

-fports *flag*

Instructs **spinFab** to include or exclude F_Ports in the testing. This feature is disabled by default. If enabled, **spinFab** tests the F_Ports connected to Brocade-branded HBAs along with other valid ports (E-ports & Loopback ports). The HBA must be running firmware v2.11 or higher. Valid flag values include the following:

0

Does not include F_Ports in the port list for testing (default).

1

Includes F_Ports in the port list for testing.

-nframes *number*

Determines how many frames will be sent by spinfab on the link to spin. Default is 2. The maximum number is 5. If you want to configure more than 5 use **-nframes 0**.

-pattern *number/name*

Forces spinfab to use a particular data pattern for the frames it spins between ports. The pattern can be a user defined pattern or is selected from a set of twenty five predefined pattern types. Use the **datatypeshow** command to view the predefined patterns supported with spinfab. For each pattern, the **datatypeshow** command displays the name, the pattern type number, and an example. Specify the pattern by its name after the **-pattern** option. If pattern is not specified, it defaults to jCRPAT (type=17). To use a user defined pattern, follow the **-pattern** option with a 32 bit hexadecimal number.

-timeout *length*

Defines a time limit (in seconds) for the running of spinfab. Spinfab will keep track of how long the test has run and stop testing once the timeout limit has been reached. The overall accuracy will be in the tens of seconds.

Diagnostics

When it detects failures, the test may report one or more of the following error messages. If errors persist, contact Technical Support.

```

DATA
ERR_STAT
ERR_STATS
ERR_STATS_2LONG
ERR_STATS_BADEOF
ERR_STATS_BADOS
ERR_STATS_C3DISC
ERR_STATS_CRC
ERR_STATS_ENCIN
ERR_STATS_ENCOUT
ERR_STATS_TRUNC
ERR_STAT_2LONG
ERR_STAT_BADEOF
ERR_STAT_BADOS
ERR_STAT_C3DISC
ERR_STAT_CRC
ERR_STAT_ENCIN
ERR_STAT_ENCOUT
ERR_STAT_TRUNC
FINISH_MSG_ERR
INIT
MBUF_STATE_ERR
NO_SEGMENT
PORT_ABSENT
PORT_DIED
PORT_ENABLE
PORT_M2M
PORT_STOPPED
PORT_WRONG
RXQ_RAM_PERR
STATS
STATS_C3FRX
STATS_FRX
STATS_FTX
TIMEOUT
XMIT

```

Examples

To test cascading ISLs:

```

switch:admin> spinfab -ports 1/0 - 1/2
spinfab running...
spinfab: Completed 11 megs, status: passed.
    port 0 test status: 0x00000000 -- passed.
    port 1 test status: 0x00000000 -- passed.
    port 2 test status: 0x00000000 -- passed.
Test Complete: "spinfab" Pass 10 of 10
Duration 0 hr, 0 min & 41 sec (0:0:41:877).
passed.

```

SEE ALSO

[itemList](#), [portLoopbackTest](#), [portPerfShow](#)

sshUtil

Manages public key authentication.

Synopsis `sshutil allowuser user name`
`sshutil showuser`
`sshutil importpubkey`
`sshutil showpubkeys`
`sshutil delpubkeys`
`sshutil genkey [-rsa | -dsa | -ecdsa]`
`sshutil exportpubkey`
`sshutil delprivkey`
`sshutil delknownhost [-all]`
`sshutil genhostkey [-rsa | -dsa | -ecdsa]`
`sshutil showhostkey`
`sshutil deletehostkey [-rsa | -dsa | -ecdsa]`
`sshutil help`

Description Use this command to enable and manage SSH public key authentication on a switch. SSH public key authentication provides a mechanism for authenticating an authorized user without a password. SSH public key authentication is more secure than password authentication and can be used to securely access services that require automatic login.

SSH public key authentication works as follows: An authorized user generates a pair of encryption keys (public and private) on a local machine (a switch or a server). Messages encrypted with the private key can only be decrypted by the public key, and vice versa. The private key remains on the local machine; the public key is exported to a remote host. The remote host responds to login requests by sending a brief message encrypted with the public key. The private key on the local host decrypts the message, and the login succeeds.

Use the **sshUtil** command to do the following:

- Configure a user to manage keys on a switch.
- Generate a private/public key pair on the local switch.
- Import a public key for a specified user from a remote host to the local switch.
- Export the public key from the local switch to a remote host.
- Delete the public keys associated with a specified user or all users on the local switch.
- Delete the private key on the local switch.
- Delete the known host name or IP address from the file `.ssh/known_hosts`.
- Generate, display, and delete the SSH host keys.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Outgoing public key authentication from the switch to a remote host is restricted to Fabric OS Commands which use secure copy (SCP), such as **configDownload** or **configUpload**.

This command supports generation of a public/private key pair on the switch to enable outgoing connections between a switch and a remote host. To set up incoming connections, you must first generate the public/private key pair on a remote host and then import the public key to the switch. Use the SSH utility **ssh-keygen -t dsa** to generate the keys on the remote host. Refer to your UNIX system documentation for details on this command.

Operands This command supports the following operands:

allowuser *user name*

Configures the specified user to perform key management operations such as to generate a key pair, export the public key, and delete the private key. The default admin is, by default, a configured user. Only one user can be configured at any given time.

The following operand is required:

user name

Specifies login name for the configured user. The user must be in the switch user database and must have admin privileges on the switch.

showuser

Displays the currently configured user. This operation can only be performed by the default admin.

importpubkey

Imports a public key from a remote host to the local switch. The protocol used is SCP. This operation supports authentication for incoming connections. For this operation to succeed, a public/private key pair must be generated on the remote host prior to the import by issuing **ssh-genkey -t dsa** (a UNIX command). The command prompts for a user name for whom the public key is imported. Once the public key is imported successfully, the user for whom the key was imported can perform public key authentication with the switch from the remote host, on which the private key resides.

importpubkey prompts for the following input parameters:

Username

Enter the user name for whom the key is imported.

IP Address

Enter the IP address for the remote host. IPv4 and IPv6 addresses are supported.

remote directory

Enter the path where the public key is stored on the remote host. The default directory where SSH stores public keys is `~username/.ssh`.

public key name

Enter the name of the file in which the public key is stored on the remote host. This is a user-generated file name that must have a `.pub` extension.

login name

Enter the user login name for the remote host.

password

Enter the password for the remote user.

showpubkeys

Displays all imported public keys associated with the specified user. Public keys generated on the switch are not shown. This command interactively prompts for a username.

username

Enter the username for whom you want to display the public keys.

delpubkey

Deletes all imported public keys associated with a specified user on the switch or with all users. This command prompts for the user name associated with the public keys. Enter "all" to delete the public keys of all users. Deletion of a configured user's public keys effectively blocks incoming connections from this user that rely on public key authentication with the switch.

genkey [-rsa | -dsa | -ecdsa]

Generates a private or public key pair for the selected type on the local switch. This option can be performed only by a configured user. This option enables authentication for outgoing connections from the switch to a remote host. You must export the public key to a remote host to complete the setup. For incoming connections, the private or public key must first be generated on the remote host by issuing **ssh-genkey -t dsa** (a UNIX command), and then importing the public key from the remote host to the switch using the **sshutil import** command.

genkey prompts for user input on the following parameters:

passphrase

Accepts a string of arbitrary length. This operand is optional, but creating a pass phrase is strongly recommended. A strong pass phrase is 10-30 characters long, fairly complex and difficult to guess, and contains a mix of upper and lowercase letters, numbers, and nonalphanumeric characters. There is no way to recover a lost pass phrase. If the pass phrase is lost, a new key must be generated and copied to the corresponding public key to other machines.

exportpubkey

Exports the public key from the switch to a specified remote host to support outgoing connections from the switch to a remote host. This option can only be performed by a configured user. The successfully exported public key must be appended to the `authorized_keys` file on the remote host. Use the `cat ~/.ssh/outgoing.pub >> ~/.ssh/authorized_keys` command to append the file.

exportpubkey prompts for IP Address, remote directory, login name and password. Refer to **importpubkey** for a description of these parameters.

delprivkey

Deletes the private key for outgoing connection from the switch. This option can only be performed by a configured user. Deletion of a configured user's private keys effectively blocks outgoing connections initiated by this user that rely on public key authentication with a remote host.

delknownhost [-all]

Deletes the known host name or IP address from the file `.ssh/known_hosts`. This option can only be performed by the authorized user. On deletion of a known host name or IP address from the `.ssh/known_hosts` file, the next SSH connection prompts the user to accept a new public key.

-all

Deletes all the known host names or IP addresses from the file.

genhostkey [-rsa | -dsa | -ecdsa]

Generates a host key and installs it on the switch.

showhostkey

Displays the host keys installed on the switch.

deletehostkey [-rsa | -dsa | -ecdsa]

Deletes the host keys installed on the switch.

help

Displays the command usage.

Examples

To configure a user for public key authentication:

```
switch:admin> sshutil allowuser username
Allowed user has been successfully changed to username.
```

To display the configured user:

```
switch:admin> sshutil showuser
username
```

To set up SSH public key authentication on a switch for incoming connections:

1. Generate a private/public key pair on a remote host (accept default directory and file name):

```
username@remotehost> ssh-keygen -t dsa
Generating public/private dsa key pair.
Enter file in which to save the key \
(/users/home/username/.ssh/id_dsa):
Enter passphrase (empty for no passphrase): passphrase
Enter same passphrase again: passphrase
Your identification has been saved in \
/users/home/username/.ssh/id_dsa.
Your public key has been saved in \
/users/home/username/.ssh/id_dsa.pub.
The key fingerprint is:
3 0:9f:ae:b6:7f:7e:55:e4:b2:7a:51:f0:95:44:5c:d1
```

2. Import the public key from the remote host to the local switch:

```
switch:username> sshutil importpubkey
Enter user name for whom key is imported: username
Enter IP address: Remote host IP Address
Enter remote directory: ~username/.ssh
Enter public key name(must have .pub suffix): id_dsa.pub
Enter login name: username
Password:*****
public key is imported successfully.
```

3. Connect to switch using remote ssh client with the `-i private_key` option:

```
username@remotehost> ssh username@IP Address-i id_dsaIP Address-i id_dsaIP Address-i id_dsa
```

To display the imported public keys on a switch:

```
switch:username> sshutil showpubkeys
Enter user name whose ssh public key is to be displayed: username
user's public keys
ssh-dss AAAAB3NzaC1kc3MAAACBANXUrSJoIA0PFJtGuZVLfqvfSr\
DYPp1WuFouOmTcmuNvpTnd+yoZ
u3C/lAu930HLTmhfxeke/NWRIdj2MJS8yTf30a0u4bf9MSNB8Pt453P/+ \
7VHHxNBYsh+Z++Dv1hfcTeb
0s53bdf7jyYSUdjlk+w//sNTaz0DCs0+rimo4l2NAAAAFQDCuHKRctSH \
D8PRYu5EelyWCQKT/wAAAAIo
AMvr1ooq0JVXmXfd0VKcC7AImzFYgRa/FOxZBe4JDkCAXztFk5wnAFy \
UbyTWEoC955mkYGqZRydMrSNM
9wLCaf2DTxXxuhFuJA1REL5NGdZqRwo2Sk5HLkYQYm1w9r9vfK \
QnFH3wYsnHV2sq7+tyRlXfwe416ee
chdwWVpmjgAAAIEAqxcaE1vY4o/cBq1Py62lPaZTcfOHS3jjdKgSO \
BKPCCVeNyx4gxnmqvihtyroEWAY
dBdk4CFgyhut16a/QmdFjn6iyiNR2SGV7X9xqkjpN8H4EhIPXGxoD \
VOFY1Vdt3V3KUXVeEI+vTBI2KJd
PmmLfyeEKZqCHOlwBx+HuuZP2BnU= username@host
```

To delete all imported public keys for a single user:

```
switch:username> sshutil delpubkeys
Enter user name for whom ssh public key is to deleted \
or all for all users:username
WARNING: It deletes all the ssh public keys for user. \
Do you want to proceed(Yes or No, default is No)?yes
ssh public keys associated to username are deleted.
```

To set up SSH public key authentication on a switch for outgoing connections:

1. Generate a private/public key pair on the local switch:

```
switch:username> sshutil genkey -rsa
Enter passphrase (empty for no passphrase):passphrase
Enter same passphrase again: passphrase
Key pair generated successfully.
```

2. Export the public key to a remote host:

```
switch:username> sshutil exportpubkey
Enter IP address: remote host IP Address
Enter remote directory: ~username/.ssh
Enter login name:username
Password:*****
public key out_going.pub is exported successfully.
```

3. Append the public key to the authorized_keys file on the remote host:

```
username@remotehost> cat ~/.ssh/outgoing.pub >> \
~/.ssh/authorized_keys
```

To delete the private key on a switch:

```
switch:username> sshutil delprivkey
private key is deleted successfully.
```

To delete a known host or IP address from the file .ssh/known_hosts:

```
switch:username> sshutil delknownhost
IP Address/Hostname to be deleted:172.26.26.104
```

To delete all the known hosts or IP addresses from the file .ssh/known_hosts:

```
switch:username> sshutil delknownhost -all
This Command will delete all the known host keys.
Please Confirm with Yes(Y,y), No(N,n) [N]: Y

All known hosts are successfully deleted.
```

SEE ALSO None

statsClear

Clears port and diagnostic statistics.

Synopsis **statsclear**
 [**--slot** *slot*]
 [**-uports** *itemlist*]
 [**-bports** *itemlist*]
 [**-use_bports** *value*]

Description Use this command to clear the port and diagnostics statistics for the specified list of blade or user ports.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following are optional:

--slot *slot* Specifies the slot on which to operate. If this option is not specified, the default slot is assumed. The default slot is 0 and designed to operate on fixed-port-count products, if **-use_bports** sets with nonzero value.

-uports *itemlist* Specifies the list of user ports for which statistics are to be cleared.

-bports *itemlist* Specifies the list of blade ports for which statistics are to be cleared.

-use_bports *value* Specify a nonzero value to clear the diagnostics statistics for the blade ports specified in **-bports** clears. A value of zero (0) clears the user ports specified in **-uports**. The default value is 0.

Examples To clear port and diagnostic statistics:

```
switch:admin> statsclear -bports 1/10-1/62 -use_bports 1
```

SEE ALSO [itemList](#)

stopPortTest

Terminates the running **portTest**.

Synopsis `stopporttest [-ports itemlist]`

Description Use this command to terminate the **portTest** command. Refer to the **portTest** help page for more information.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

-ports *itemlist*

Terminates the test on the specified ports; this operand is optional; if omitted, the test is terminated on all ports. Refer to the **itemList** help page for more information.

Examples To stop the **portTest** command:

```
switch:admin> stopporttest
```

SEE ALSO [portLoopbackTest](#), [portTest](#), [portTestShow](#), [spinFab](#)

supportFfdc

Modifies or displays the first-fault data capture (FFDC) daemon.

Synopsis `supportffdc [--disable | --enable | --show]`

Description Use this command to disable or enable the FFDC events, or to display the current configuration. If disabled, the daemon does not capture any data even when a message with FFDC attributes is logged. FFDC is enabled by default. When executed without operands, the command prints the usage.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------|---|
| --disable | Disables the FFDC. |
| --enable | Enables the FFDC. |
| --show | Displays the FFDC configuration parameters. |

Examples To display the FFDC configuration:

```
switch:admin> supportffdc --show
First Failure Data Capture (FFDC) is disabled.
```

To enable the FFDC events:

```
switch:admin> supportffdc --enable
First Failure Data Capture (FFDC) is enabled.
```

SEE ALSO None

supportFtp

Sets, clears, or displays support FTP parameters and enables or disables auto file transfer.

Synopsis `supportftp [-S]`

`supportftp -s [-h host][-u username][-p password]
[-d remotedirectory] [-l protocol]`

`supportftp -t hours`

`supportftp -R`

`supportftp -e`

`supportftp -d`

Description Use this command to set, clear, or display **supportFtp** parameters. The parameters set by this command are used by the **supportSave** and **traceDump** commands.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-S

Displays the current **supportFtp** parameters.

-s

Sets the **supportFtp** parameters. The following operands are optional. If the **-s** option is specified without further operands, the command interactively prompts for these parameters.

-h *host*

Specifies the host. Provide an IP address or a server name. IPv4 and IPv6 addresses are supported. To specify the host by name, a DNS entry must exist for the server.

-u *username*

Specifies the user name. The user name must be less than 48 characters long.

-p *password*

Specifies the account password. The password must be less than 48 characters long. When using the reserved user names such as *anonymous* or *FTP*, a password is not required.

-d *remotedirectory*

Specifies the remote directory where the trace dump files are stored. The directory name must be less than 48 characters long. Specifying the root directory as the remote directory (/) is not allowed.

-l *protocol*

Specifies the transfer protocol. Valid values are file transfer protocol (FTP), secure copy protocol (SCP), or secure FTP (SFTP).

- t hours** Specifies the time interval, in units of hours, at which the server connectivity is checked. The fractional value for time interval is converted to the nearest lower integer value, for example, a value of 6.5 hours is truncated to 6 hours.
- R** Clears all **supportFtp** parameters.
- e** Enables auto file transfer. Trace dump files are automatically transferred to a designated FTP server. The server parameters must be set before you can enable auto file transfer.
- d** Disables auto file transfer.

Examples

To set the FTP parameters:

```
switch:admin> supportftp -s -h 1080::8:800:200C:417A \
-u admin -p password -d support -l sftp
supportftp: ftp parameters changed.
```

To display the FTP parameters:

```
switch:admin> supportftp
Host IP Addr: 1080::8:800:200C:417A
User name: admin
Remote Dir: support
Auto Upload protocol: sftp
Auto-FTP: Off
```

To set FTP parameters interactively:

```
switch:admin> supportftp -s
Host IP Addr[1080::8:800:200C:417A]:192.168.67.126
User Name[admin]: admin
Password[*****]:password
Remote Dir[support]:
Auto upload protocol[ftp]:scp
Auto file transfer parameters changed
```

To set the time interval at which the FTP server connectivity is checked:

```
switch:admin> supportftp -t 24
supportftp: ftp check period changed.
```

To enable auto file transfer:

```
switch:admin> supportftp -e
support auto file transfer enabled.
```

To disable auto file transfer:

```
switch:admin> supportftp -d
support auto file transfer disabled.
```

SEE ALSO [supportSave](#), [supportShow](#), [traceDump](#)

supportInfoClear

Clears all the default port statistic counters and portlogs in a chassis or switch.

Synopsis `supportinfoclear --clear [-RASlog] [-force]`

`supportinfoclear --help`

Description Use this command to clear all the default port statistic counters, and portlogs in a chassis or switch. Use the **-RASlog** option to clear error logs along with statistics and portlogs on the active CP. On the standby CP, the **-RASlog** option clears only the error logs.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--clear

Clears port statistic counters and portlogs in a chassis or switch. The command prompts for a confirmation.

-RASlog

Clears error logs along with port statistic counters and portlogs in a chassis or switch. The command prompts for a confirmation.

-force

Executes without prompting for a confirmation.

--help

Displays command usage.

Examples To clear port statistic counters and portlogs:

```
switch:admin> supportinfoclear --clear
Execution of the command will clear the default collection of counters and port logs
in the chassis.
Would you like to continue [y/n]?: y
```

To clear port statistic counters and portlogs along with error logs:

```
switch:admin> supportinfoclear --clear -RASlog
Execution of the command will clear the error logs along with default collection of
counters and port logs in the chassis.
Would you like to continue [y/n]?: y
```

SEE ALSO [errClear](#), [portLogClear](#), [portStatsClear](#), [statsClear](#)

supportSave

Saves RASLOG, TRACE, **supportShow**, core file, FFDC data, and other support information

Synopsis

supportsave

```
supportsave [-n] [-c] [-k] [-a]
                [-u user_name -p password -h host_ip
                -d remote_dir -l protocol]
```

```
supportsave [-R]
```

```
supportsave [-U -d remote_dir]
```

```
supportsave [-t timeout_multiplier]
```

Description

Use this command to collect RASLOG, TRACE, **supportShow**, core file, FFDC data and other support information to a remote FTP location. On platforms that support USB, the information can also be stored on an attached USB device. On a dual-CP system, information is saved for the local and the remote CP. **SupportShow** information is available on Active and Standby CPs. To reduce the chance of missing the correct trace dump, **supportSave** retrieves old (the dump created prior to the current one) and new (the dump triggered by the command) trace dumps.

The files generated by this command are compressed before being sent off the switch. The core files and panic dumps remain on the switch after the command is run. The FFDC data are removed after the command has finished.

If there are blade processor (BP) blades installed on the switch, a support file (a.tar.gz file) is generated from each slot.

This command accepts IPv4 and IPv6 addresses. If the configured IP address is in IPv6 format, the RAS auto file transfer and event notification to syslog will not work in the case where the Fabric OS version is downgraded. You must reconfigure auto file transfer and syslog with IPv4 IP addresses.

In a Virtual Fabric environment, **supportSave** saves all chassis-based information and iterates through the defined switch-based information for all logical switches. Chassis permissions are required to execute this command.

Note that quotes should be used around path entries to ensure proper handling of special shell characters.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

If you use anonymous FTP to run **supportSave** on a chassis with multiple AP blades, configure the FTP Windows server to allow unlimited anonymous users.

Operands

When invoked without operands, this command goes into interactive mode. The following operands are optional:

-n

Does not prompt for confirmation. This operand is optional; if omitted, you are prompted for confirmation.

supportSave

- c** Uses the FTP, SCP, or SFTP parameters saved by the **supportFtp** command. This operand is optional; if omitted, specify the FTP, SCP, or SFTP parameters through command line options or interactively. To display the current **supportFTP** parameters, run **supportFtp** (on a dual-CP system, run **supportFtp** on the active CP).
The **-c** option is mutually exclusive with **-u**, **-p**, **-h**, and **-d**.
- k** Used to transfer the core and FFDC files. This operand can be specified with the **-c** and **-n** operands.
- a** Enables Challenge Response Authentication (CRA). CRA is supported only with the SCP protocol.
- u *user_name*** Specifies the user name for the FTP, SCP, or SFTP server. This operand is optional; if omitted, anonymous FTP is used.
- p *password*** Specifies the password for the FTP, SCP, or SFTP server. This operand is optional with FTP; if omitted, anonymous FTP is used.
- h *host_ip*** Specifies the IPv4 or IPv6 address for the remote server.
- d *remote_dir*** Specifies the remote directory to which the file is to be transferred. By default, the remote directory is created in the **/support** directory of the USB device.
- R** Removes all core files on the CP and BP. This option cannot be used with any other options.
- l *protocol*** Specifies the transfer protocol. Valid values are file transfer protocol (FTP), secure copy (SCP), or secure FTP (SFTP).

If you plan to use SCP to transfer files, it is important to test the **supportSave** command prior to its use with various SCP-mode services. Because the **supportSave** command makes several access requests to copy files, it is important that the SCP-mode service be configured so that passwords are not required for each attempted transfer by the **supportSave** command. Failure to configure the service correctly may result in significant delays in obtaining transferred output from the **supportSave** command.

When using secure copy (SCP), **supportSave** may create a directory specified by the **-d** option if it does not already exist and the parent directory has the appropriate permissions. Use of FTP requires the directory to exist on the remote server.
- U** Saves support data to an attached USB device. When using this option, a target directory must be specified with the **-d** option.

-t timeout_multiplier

Extends predefined **SupportSave** timeout values by the value of the timeout multiplier. Use this option to repeat the **supportSave** operation when **supportSave** completion indicates that one or more modules timed out during the process. For example, **-t 2** doubles the timeout values for each of the **SupportSave** modules. Valid multiplier values are 2 to 5. The default is 1.

Examples

To save RASLOG, TRACE, **supportShow**, and other support information to an FTP server in interactive mode:

```
switch:admin> supportsave
This command collects RASLOG, TRACE, supportShow, \
  core file, FFDC data
and then transfer them to a FTP/SCP/SFTP server \
  or a USB device.
This operation can take several minutes.
NOTE: supportSave will transfer existing trace dump \
  file first, then
  automatically generate and transfer latest one. \
  There will be two trace dump
  files transferred after this command.
OK to proceed? (yes, y, no, n): [no] y

Host IP or Host Name: 192.168.126.115
User Name: admin
Password:*****
Network Protocol(1-auto-select, 2-FTP, 3-SCP, 4-SFTP) [1]: 4
Remote Directory: /temp/support

Do you want to continue with CRA (Y/N) [N]: Y

Saving support information for chassis:HL51,module:RAS...
Saving support information for chassis:HL51,module:TRACE_OLD...
Saving support information for chassis:HL51,module:TRACE_NEW...
Saving support information for chassis:HL51,module:FABRIC...
Saving support information for chassis:HL51,module:CORE_FFDC...
Saving support information for chassis:HL51,module:DIAG...
Saving support information for chassis:HL51,module:RTE...
Saving support information for chassis:HL51,module:ISCSID_DBG...
Saving support information for chassis:HL51,module:AGDUMP...
Saving support information for chassis:HL51,module:SSHOW_PLOG...
Saving support information for chassis:HL51,module:SSHOW_OS...
Saving support information for chassis:HL51,module:SSHOW_EX...
Saving support information for chassis:HL51,module:SSHOW_FABRIC...
(output truncated)
```

To collect support information on a Brocade 5100 and save it to an attached USB device (timeout values are doubled):

```
switch:admin> supportsave -U -d mysupportsave -t 2
This command collects RASLOG, TRACE, supportShow, \
  core file, FFDC data
and then transfer them to a FTP/SCP/SFTP server \
  or a USB device.
This operation can take several minutes.
NOTE: supportSave will transfer existing trace dump \
  file first, then
  automatically generate and transfer latest one. \
  There will be two trace dump
```

supportSave

```
files transferred after this command.
OK to proceed? (yes, y, no, n): [no] y

Saving support information for chassis:ras9,module:RAS...
Saving support information for chassis:ras9,module:TRACE_OLD...
Saving support information for chassis:ras9,module:TRACE_NEW...
Saving support information for chassis:ras9,module:FABRIC...
Saving support information for chassis:ras9,module:CORE_FFDC...
No core or FFDC data files found!
Saving support information for chassis:ras9,module:DIAG..
Saving support information for chassis:ras9,module:RTE...
Saving support information for chassis:ras9,module:ISCSID_DBG...
Saving support information for chassis:ras9,module:AGDUMP...
Saving support information for chassis:ras9,module:SSHOW_PLOG...
(output truncated)
```

To run **supportSave** without confirmation on a Brocade DCX with AP blades included using **supportFTP** parameters (only Active CP output is shown):

```
switch:admin> supportsave -n -c
Saving support information for chassis:ras2,module:RAS...
Saving support information for chassis:ras2,module:TRACE_OLD...
Saving support information for chassis:ras2,module:TRACE_NEW...
Saving support information for chassis:ras2,module:FABRIC...
Saving support information for chassis:ras2,module:CORE_FFDC...
Saving support information for chassis:ras2,slot:4...
slot 4 support file transfer done.
Saving support information for chassis:ras2,slot:12...
slot 12 support file transfer done.
Saving support information for chassis:ras2,module:DIAG...
Saving support information for chassis:ras2,module:RTE...
Saving support information for chassis:ras2,module:ISCSID_DBG...
Saving support information for chassis:ras2,module:AGDUMP...
Saving support information for chassis:ras2,module:SSHOW_PLOG...
(output truncated)
```

SEE ALSO [supportShow](#), [supportFtp](#)

supportShow

Displays switch information for debugging purposes.

Synopsis **supportshow**

supportshow *slot[/port1-port2] [lines]*

Description

Use this command to display support information from groups of preselected Fabric OS and Linux commands and other support and debugging information. You can specify a port or a range of ports for which to display this information. These commands are organized by groups, but note that the order of the groups listed below is not the same as executed by the command.

The FCIP commands are supported only on the Brocade 7800 switch and the Brocade FX8-24 blade. On unsupported platforms, the command displays a "not applicable to this platform" message next to the FCIP command group.

SupportShow executes commands in the following command groups. Use **supportShowCfgenable** or **supportShowCfgDisable** to modify the settings for each group.

| | |
|------------------|---|
| os | OS group commands, enabled by default. |
| exception | Exception group commands, enabled by default. |
| port | Port group commands, enabled by default. |
| fabric | Fabric group commands, enabled by default. |
| services | Service group commands, enabled by default. |
| security | Security group commands, enabled by default. |
| network | Network group commands, enabled by default. |
| portlog | Portlog group commands, enabled by default. |
| system | System group commands, enabled by default. |
| extend | Extend group commands, disabled by default. |
| filter | Filter group commands, disabled by default. |
| ficon | FICON group commands, disabled by default. |

| | |
|----------------|---|
| iswitch | FC Router group commands, disabled by default. |
| asic_db | ASIC DB group commands, disabled by default. |
| fcip | FCIP group commands, disabled by default. |
| ag | Access Gateway group commands, disabled by default. |
| dce_hsl | DCE group commands, enabled by default. |
| crypto | Encryption group commands, disabled by default. |

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This is a diagnostic command and should only be run for diagnostic support.

Output generated by this command may vary by switch configuration and platform. Output may change without notice.

Operands This command has the following operands:

| | |
|--------------------|--|
| slot | On bladed systems only, specifies a slot number. |
| port1-port2 | Specifies a port or a range of ports for which to display supportShow information. This operand is optional; if omitted, the command displays information for all ports. The following values are supported: <ul style="list-style-type: none"> • A single port in <i>port1-port1</i> format, for example, "5/8-8". • A port range where the beginning and end port are separated by a dash, for example, "5/8-13". |
| lines | Specifies the number of lines for the portLogDump output. This parameter is valid only with the <i>slot/port</i> parameters. |

Examples To display debugging information for a single port on a Brocade 5300:

```
switch:admin> supportshow 43
VF
=====
Date:
Sun Dec 6 05:10:13 PST 2009

Time Zone:
America/Los_Angeles

Version:
Kernel:      2.6.14.2
```

```
Fabric OS: v6.4.0_main_bld09
Made on:   Tue Dec 1 20:04:36 2009
Flash:    Wed Dec 2 11:54:49 2009
BootProm: 1.0.15
```

```
supportshow groups enabled:
  os          enabled
  exception   enabled
  port        enabled
  fabric      enabled
  services    enabled
  security    enabled
  network     enabled
  portlog     enabled
  system      enabled
  extend      disabled
  filter      disabled
  ficon       disabled
  iswitch     enabled
  asic_db     enabled
  fcip        disabled (not applicable to this platform)
  ag          enabled
  dce_hsl     enabled
```

```
**** Begin start_port_log_cmd group ****
Sun Dec 6 05:10:14 PST 2009
portlogdump:
CURRENT CONTEXT -- 0 , 128
portlogdump      :
time            task      event  port cmd  args
-----
Sat Dec 5 23:54:37 2009
23:54:37.560 FCPH      read   56 16 02fffc23,00fffc19, \
      bb000000,00000000,04 3401bb
23:54:37.560 FCPH      seq    56 10 20290000,043401bb, \
      00000722,0000001c,00 000000
23:54:37.560 msd0      ctin   56 fa 0001f007,00000000
23:54:37.561 msd0      ctout  56 fa 00018001,0009f300
(output truncated)
```

SEE ALSO [supportFtp](#), [supportSave](#), [supportShowCfgDisable](#), [supportShowCfgEnable](#), [supportShowCfgShow](#), [traceDump](#)

supportShowCfgDisable

Disables a group of commands under the **supportShow** command.

Synopsis `supportshowcfgdisable os | exception | port | fabric
| services | security | network | portlog | system | extend
| filter | ficon | iswitch | asic_db | ag
| dce_hsl | crypto | fcip`

Description Use this command to disable a group of commands under the **supportShow** command. Use the **supportShowCfgEnable** command to enable groups of commands.

The FCIP commands are supported only on the Brocade FX8-24 platform and cannot be configured to collect data on unsupported platforms.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------|--|
| os | Disables the OS group commands. |
| exception | Disables the exception group commands. |
| port | Disables the port group commands. |
| fabric | Disables the fabric group commands. |
| services | Disables the service group commands. |
| security | Disables the security group commands. |
| network | Disables the network group commands. |
| portlog | Disables the portlog group commands. |
| system | Disables the system group commands. |
| extend | Disables the extend group commands. |
| filter | Disables the filter group commands. |
| ficon | Disables the FICON group commands. |

| | |
|----------------|---|
| iswitch | Disables the FC Router group commands. |
| asic_db | Disables the ASIC DB group commands. |
| ag | Disables the Access Gateway group commands. |
| dce_hsl | Disables the DCE group commands. |
| crypto | Disables the encryption group commands. |
| fcip | Disables the FCIP group commands. Supported only on the Brocade FX8-24. |

Examples To disable the OS group of commands under the **supportShow** command:

```
switch:admin> supportshowcfgdisable os  
Config update Succeeded
```

SEE ALSO [supportShow](#), [supportShowCfgEnable](#), [supportShowCfgShow](#)

supportShowCfgEnable

Enables a group of commands to be displayed under the **supportShow** command.

Synopsis `supportshowcfgenable os | exception | port | fabric
| services | security | network | portlog | system
| extend | filter | ficon | iswitch | asic_db
| ag | dce_hsl | crypto | fcip`

Description Use this command to enable a group of commands to be displayed under the **supportShow** command. Use the **supportShowCfgDisable** command to disable groups of commands.

The FCIP commands are supported only on the Brocade FX8-24 platform and cannot be configured to collect data on unsupported platforms.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

| | |
|------------------|---------------------------------------|
| os | Enables the OS group commands. |
| exception | Enables the exception group commands. |
| port | Enables the port group commands. |
| fabric | Enables the fabric group commands. |
| services | Enables the service group commands. |
| security | Enables the security group commands. |
| network | Enables the network group commands. |
| portlog | Enables the portlog group commands. |
| system | Enables the system group commands. |
| extend | Enables the extend group commands. |
| filter | Enables the filter group commands. |
| ficon | Enables the FICON group commands. |

| | |
|----------------|---|
| iswitch | Enables the FC Router group commands. |
| asic_db | Enables the ASIC DB group commands. |
| ag | Enables the Access Gateway group commands. |
| dce_hsl | Disables the DCE group commands |
| crypto | Enables the encryption group commands. |
| fcip | Enables the FCIP group commands. Supported only on the Brocade FX8-24 platform. |

Examples To enable a group of commands under the **supportShow** command:

```
switch:admin> supportshowcfgenable os
Config update Succeeded
```

SEE ALSO [supportShow](#), [supportShowCfgDisable](#), [supportShowCfgShow](#)

supportShowCfgShow

Displays the groups of commands enabled for display by the **supportShow** command.

Synopsis `supportshowcfgshow`

Description Use this command to display the groups of commands enabled for display by the **supportShow** command. Use the **supportShowCfgEnable** and the **supportShowCfgDisable** commands to modify which groups are displayed.

The FCIP commands are supported only on the Brocade FX8-24 platform and cannot be configured to collect or display data on unsupported platforms.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display command groups configured for data collection on the Brocade 6510:

```
switch:admin> supportshowcfgshow
os          enabled
exception   enabled
port        enabled
fabric      enabled
services    enabled
security    enabled
network     enabled
portlog     enabled
system      enabled
extend      disabled
filter      disabled
ficon       disabled
iswitch     disabled
asic_db     enabled
fcip        disabled (not applicable to this platform)
ag          disabled
```

To display command groups configured for data collection on the Brocade 7800

```
switch:admin> supportshowcfgshow
os          enabled
exception   enabled
port        enabled
fabric      enabled
services    enabled
security    enabled
network     enabled
portlog     enabled
system      enabled
extend      disabled
filter      disabled
ficon       disabled
```

```
iswitch    enabled
asic_db    enabled
fcip       enabled
ag         enabled
dce_hsl    enabled
```

SEE ALSO [supportShow](#), [supportShowCfgDisable](#), [supportShowCfgEnable](#)

switchBeacon

Sets switch beaconing mode on or off.

Synopsis `switchbeacon [mode]`

Description Use this command to enable or disable switch beaconing mode. Switch beaconing can be used to locate a failing unit.

When beaconing mode is turned on, the port LEDs flash amber, left to right and right to left, from port 0 to the highest port number and back to port 0. The beaconing mode continues until you turn it off.

The beaconing LED pattern continues until you turn it off. Beaconing mode takes over the port LEDs. Other commands are still executable and functional. The normal flashing LED pattern (associated with an active, faulty or disabled port for example) is suppressed and only the beaconing pattern is shown. However, if diagnostic frame-based tests (such as **portLoopbackTest**) are executed, two patterns are interleaved. The diagnostic test flickers the LEDs green and the beaconing mode runs the LEDs amber at the same time.

The **switchBeacon** command is one of the commands that controls beaconing. Each command has a clearly defined scope of action:

- The **portBeacon** command enables or disables beaconing on a specified port.
- The **switchBeacon** command enables or disables beaconing on all ports in the current logical switch.
- The **chassisBeacon** command enables or disables beaconing on all ports in the chassis.
- The **portPeerBeacon** command enables or disables beaconing to identify the interconnections between ports.

The actions of the beaconing commands are independent and mutually exclusive. For example, if you enabled beaconing on the logical switch (1) and you want to enable beaconing on the entire chassis, you must first disable switch beaconing with the **switchBeacon** command before you can use the **chassisBeacon** command to enable beaconing on the entire chassis. Likewise, existing **portBeacon** settings remain unaffected if you enable or disable beaconing on the switch or on the chassis. Failure to disable existing beaconing commands before using a different type of beaconing may cause the commands to interfere with each other in unexpected ways.

To determine beaconing status, use the **switchBeacon** or **chassisBeacon** command without operands. A value of 0 indicates that the command is disabled, a value of 1 indicates that the command is enabled. Issue the **portBeacon --show** command to display beaconing for a specific port. The **switchShow** command displays the status of the **switchBeacon** command only.

Notes The **switchBeacon** command is not supported on the Brocade FCOE10-24 blade.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

mode

Specify 1 to enable beaconing mode or 0 to disable beaconing mode. This operand is optional. If no operand is specified, the current value is displayed.

Examples To turn beaconing mode on and to verify the configuration:

```
switch:admin> switchbeacon 1  
switch:admin> switchbeacon  
value = 1
```

To turn beaconing mode off to verify the configuration::

```
switch:admin> switchbeacon 0  
switch:admin> switchbeacon  
value = 0
```

SEE ALSO [chassisBeacon](#), [portBeacon](#), [portPeerBeacon](#), [switchShow](#)

switchCfgPersistentDisable

Disables a switch persistently.

Synopsis **switchcfgpersistentdisable**
switchcfgpersistentdisable --setdisablestate
switchcfgpersistentdisable --disable
switchcfgpersistentdisable --help

Description Use this command to persistently disable the switch immediately or after reboot. All Fibre Channel ports are taken offline. If the switch was part of a fabric, the remaining switches reconfigure. The switch remains disabled even after a reboot.

As each port is disabled, the front panel LEDs change to a slow-flashing amber.

You can temporarily enable a persistently disabled switch with the **switchEnable** command. A temporarily enabled switch remains disabled after a reboot.

Notes Performance Monitoring cannot be added to any port on a persistently disabled switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--setdisablestate Sets the switch in disable state. The switch is disabled persistently after the next reboot.

--disable Disables the switch persistently.

--help Displays command usage.

Examples To disable a switch persistently:

```
switch:admin> switchcfgpersistentdisable
Switch's persistent state set to 'disabled'
```

To set the state of a switch to disable so that the switch is disabled during next reboot:

```
switch:admin> switchcfgpersistentdisable -setdisablestate
Switch's persistent state set to 'disabled'
Switch persistent disable set
```

SEE ALSO [switchDisable](#), [switchEnable](#), [switchCfgPersistentEnable](#), [switchShow](#)

switchCfgPersistentEnable

Enables a switch persistently.

Synopsis `switchcfgpersistentenable`

Description Use this command to persistently enable a persistently disabled switch. All Fibre Channel ports that passed the power-on self-test (POST) are enabled and come online if connected to a device, or remain offline if disconnected. The switch may need to be enabled if it was previously disabled to make configuration changes or to run diagnostics.

If the switch is connected to a fabric, it rejoins the fabric. If this switch remains the principal switch, it assigns itself a domain ID. If another switch assumes the principal role, then this switch becomes a subordinate switch, and accepts a domain ID from the principal. Refer to the FC-SW specification for a complete description of this process.

As each port is enabled, the front panel LEDs change from slow-flashing amber to green for online ports, or to nonflashing amber for ports that do not initialize. Disconnected ports remain unlit.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To persistently enable a previously persistently disabled switch:

```
switch:admin> switchcfgpersistentenable  
Switch's persistent state set to 'enabled'
```

SEE ALSO [switchDisable](#), [switchEnable](#), [switchCfgPersistentDisable](#), [switchShow](#)

switchCfgSpeed

Configures the speed for all ports on a switch.

Synopsis `switchcfgspeed speed`

Description Use this command to configure the port speed on a switch. This command sets the speed for all user ports. If any port on the switch is not capable of the specified speed setting, an error message is displayed for that port. The configuration is saved in nonvolatile memory and persists across switch reboots or power cycles.

On 16 Gbps-capable platforms, the **switchCfgSpeed** command checks if the requested speed is allowed, based on the combination configured for the octet that contains the first eight physical ports of the switch or blade. If the requested speed is not supported by the current octet speed combination, this command logs a RASlog message and moves on to the next port. Use the **portCfgOctetSpeedCombo** command to set the suggested combination before re-executing the **switchCfgSpeed** command.

Use the **portShow** command to display actual port speed settings. Use the **portCfgShow** command to display user-configured speed settings.

Notes This configuration cannot be set on VE_Ports or VEX_Ports.

Speed configuration is not applicable to FCoE ports.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

| | |
|--------------|---|
| <i>speed</i> | Specifies the speed for all ports on a switch. This operand is required. Valid values are as follows: |
| 0 | Autosensing mode. The port automatically configures for the highest speed. |
| 1 | The port is set at a fixed speed of 1 Gbps (not supported on 10G and 16G platforms). |
| 2 | The port is set at a fixed speed of 2 Gbps. |
| 4 | The port is set at a fixed speed of 4 Gbps. |
| 8 | The port is set at a fixed speed of 8 Gbps. |
| 16 | The port is set at a fixed speed of 16 Gbps (not supported on 8G platforms). |
| 32 | The port is set at a fixed speed of 32 Gbps (not supported on 8G platforms). |

Examples To set the autosensing mode for all ports on a switch:

```
switch:admin> switchcfgspeed 0  
Committing configuration...done.
```

SEE ALSO [portCfgOctetSpeedCombo](#), [portCfgSpeed](#), [portShow](#)

switchCfgTrunk

Enables or disables trunking on all the ports of a switch.

Synopsis `switchcfgtrunk mode`

Description Use this command to enable or disable trunking on all the ports of a switch. Use `portCfgTrunkPort` to enable or disable trunking on a single port.

When the command is executed to update the trunking configuration, the ports to which the configuration applies are disabled and subsequently re-enabled with the new trunking configuration. Traffic through these ports may be temporarily disrupted. The command issues a message that lists the VE/VEX_Ports to which the configuration does not apply.

Although trunking configuration changes are applied at the switch level, they are tracked as a per-port attribute and no switch-wide attribute is maintained to keep track of these changes. Whenever a new port comes online as part of the switch, you must reapply the trunking configuration. For example, If you remove a blade from a chassis while trunking is enabled for the ports on the blade, and you disable trunking on the switch after you removed the blade, the blade ports will come online with trunking enabled after you reinsert the blade. To avoid potentially disruptive behavior, reapply the trunking configuration.

Trunking on Inter-Chassis Link (ICL) ports is always enabled and cannot be turned off by this command.

Disabling trunking fails if a Trunk Area (TA) is enabled on the port. Use the `portTrunkArea` command to disable the TA on all ports before disabling trunking.

Notes Enabling trunking requires an ISL Trunking license. You may disable trunking without a license.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operand is required:

mode
Specify 1 to enable trunking on all ports. Specify 0 to disable trunking on all ports.

Examples To enable trunking on all ports of a switch:

```
switch:admin> switchcfgtrunk 1
Configuration applied to all ports except the following \
VE/VEX_Ports (ports 176 - 191).
```

To disable trunking on all ports of a switch:

```
switch:admin> switchcfgtrunk 0
Committing configuration...done.
```

SEE ALSO [portCfgShow](#), [portCfgTrunkPort](#), [portShow](#), [portTrunkArea](#), [switchShow](#)

switchDisable

Disables all user ports on a switch.

Synopsis **switchdisable**

Description Use this command to disable all user ports on a switch. All Fibre Channel ports are taken offline. If the switch was part of a fabric, the remaining switches reconfigure. As each port is disabled, the front panel LED changes to a slow-flashing amber.

The switch must be disabled before making configuration changes or before running offline diagnostic tests. Commands that require the switch to be disabled generate an error message if invoked while the switch is enabled. It is not necessary to disable the switch before rebooting or powering off.

When this command is executed on a logical switch, only the ports allocated to the logical are disabled. To disable the entire chassis, use the **chassisDisable** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To disable the switch:

```
switch:admin> switchdisable
```

SEE ALSO [bladeEnable](#), [bladeDisable](#), [chassisEnable](#), [chassisDisable](#), [switchCfgPersistentDisable](#), [switchCfgPersistentEnable](#), [switchEnable](#), [switchShow](#)

switchEnable

Enables all user ports on a switch.

Synopsis **switchenable**

Description Use this command to enable all user ports on a switch. All Fibre Channel ports that passed the power-on self test (POST) are enabled. They can come online if connected to a device, or remain offline if disconnected. Use **switchEnable** to re-enable the switch after making configuration changes or running offline diagnostics.

If the switch is connected to a fabric, it rejoins the fabric. If the switch remains the principal switch, it assigns itself a domain ID. If another switch assumes the principal role, then the re-enabled switch becomes a subordinate switch and accepts a domain ID from the principal.

As each port is enabled, the front panel LED changes to green for online ports, or to amber for uninitialized ports. Disconnected ports remain unlit.

When this command is executed on a logical switch, only the ports allocated to the logical switch are enabled. To enable the entire chassis, use the **chassisEnable** command.

Notes This command also enables the ports of a disabled blade on the switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To enable a switch:

```
switch:admin> switchenable
```

SEE ALSO [bladeEnable](#), [bladeDisable](#), [chassisDisable](#), [chassisEnable](#), [switchCfgPersistentDisable](#), [switchCfgPersistentEnable](#), [switchDisable](#), [switchShow](#)

switchName

Displays or sets the switch name.

Synopsis `switchname [name]`

Description Use this command to display or set the switch name. Once you set the switchname, you must re-login for the change to be in effect. All switches have a symbolic name that is primarily used for switch management. This name is shown in the Fabric OS CLI prompt, under each switch icon in Web Tools, in all the switch event RASLog messages, and in the output of various Fabric OS Commands, such as **fabricShow**.

Use this command with the *name* operand to assign a new switch name. The switch name should be unique. Enter this command without an operand to display the current switch name.

Changing the switch name causes a domain address format registered state change notification (RSCN) to be issued. Refer to the FC-FLA specification for a description of RSCNs).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

name

Specifies the name for the switch. A switch name can include up to 30 characters. The name can either begin with a letter or number and can consist of letters, numbers, hyphens, and underscore characters. The switchname which begins with numeric character should have atleast one alphabet. Switchname with all numeric character is not valid. Example switchname 12345678 is invalid. The switch name should be unique. Spaces are not allowed. When FICON Management Server (FMS) mode is enabled, the switch name can include up to 24 characters only and if the switch name has a dash (-), it is converted to a period (.). For example, the switch name "123-switch" is converted to "123.switch". This operand is optional; if omitted, this command displays the current switch name.

Examples To change the switch name to a name starting with a character (note the change in the prompt text):

```
switch:admin> switchname brocade_demo_1298765_AY4TYI60
Committing configuration...
Done.
Switch name has been changed.Please re-login into the switch for the change to be
applied.
```

switchName

To change the switch name to a name starting with a character (note that the dash (-) is converted to period (.) in the switch name because FMS mode is enabled):

```
switch:admin> switchname brocade-174
Committing configuration...
Done.
Switch name has been changed.Please re-login into the switch for the change to be
applied.
2015/11/30-13:13:43, [IPAD-1002], 191, FID 128, INFO, brocade-174, Switch name has
been successfully changed to brocade-174.

brocade-174:admin> Committing configuration...
Done.
2015/11/30-13:13:45, [IPAD-1002], 192, FID 128, INFO, brocade.174, Switch name has
been successfully changed to brocade.174.
brocade.174:admin> switchname
brocade.174
```

To change the switch name to a name starting with number (FMS mode is enabled):

```
switch:admin> switchname 174-brocade
Committing configuration...
Done.
Switch name has been changed.Please re-login into the switch for the change to be
applied.
2015/11/30-13:13:43, [IPAD-1002], 191, FID 128, INFO, 174-brocade, Switch name has
been successfully changed to 174-brocade.

174-brocade:admin> Committing configuration...
Done.
2015/11/30-13:13:45, [IPAD-1002], 192, FID 128, INFO, 174.brocade, Switch name has
been successfully changed to 174.brocade.
174.brocade:admin> switchname
174.brocade
```

To change the switch name to a name starting with number (FMS mode is disabled):

```
switch:admin> switchname 174-brocade
Committing configuration...
Done.
Switch name has been changed.Please re-login into the switch for the change to be
applied.
2015/11/30-13:13:43, [IPAD-1002], 191, FID 128, INFO, 174-brocade, Switch name has
been successfully changed to 174-brocade.
174-brocade:admin> switchname
174-brocade
```

SEE ALSO [chassisShow](#), [switchShow](#)

switchShow

Displays switch and port status.

Synopsis

switchshow

switchshow [-slot *s/of*] [-portname | -qsfp]

switchshow [-portcount]

Description

Use this command to display switch, blade, and port status information. Output may vary depending on the switch model.

When used without operands, **switchShow** displays the following information:

| | |
|-------------------------------|--|
| switchName | Switch name. |
| switchType | Switch model and revision numbers. |
| switchState | Switch state: Online, Offline, Testing, or Faulty. When you issue the switchcfgPersistentDisable command followed by the switchEnable command, the switch state changes to, "Online (Temporary)". The switch remains in this state until you issue the switchcfgPersistentEnable command. |
| switchMode | Switch operation mode: Access Gateway (if AG is enabled). |
| switchRole | Switch role: Principal, Subordinate, or Disabled. |
| switchDomain | Switch domain ID: 1 to 239. |
| switchId | Switch embedded port D_ID. |
| switchWwn | Switch world wide name (WWN). |
| switchBeacon | Switch beaconing state: On or Off. |
| zoning | The name of the active zone is displayed in parentheses. Active only when Access Gateway mode is disabled. |
| FC Router | FC Router state: On or Off. |
| FC Router BB Fabric ID | The backbone fabric ID for FC routing. |

| | |
|--|--|
| Fabric Name | The name assigned to the fabric. The fabric name is set with the fabricName command. |
| HIF Mode | Indicates HIF mode of the switch. Values are ON or OFF. The HIF mode can be set using the configure command. |
| Allow XISL Use | Allows the switch to use interswitch links (XISL) in the base fabric to carry traffic to this logical switch. Values are ON or OFF. |
| LS Attributes | On a switch in Virtual Fabric mode, this field displays logical switch attributes, including the fabric ID associated with the logical switch, the switch role (default switch or base switch), and the fabric Address Mode (0, 2 or 3). If Virtual Fabrics are disabled, only the Address Mode is displayed. The fabric Address Mode value is set by the configure command (Enable a 256 Area Limit). |
| The switch summary is followed by one-line description for non-EX_Ports and one or two lines for EX_Ports: | |
| Index | Port index is a number between 0 and the maximum number of supported ports on the platform. The port index identifies the port number relative to the switch. |
| Slot | Slot number; 1-12. |
| Port | Port number; 0-15, 0-31, or 0-63. |
| PortWWN | Port world wide name (WWN). |
| Address | The 24-bit Address Identifier. |
| Media | Media types include the following: |
| -- | No module present. Applicable to all port types. |
| cu | Displays when the copper (default) GbE port ge0 or ge1 is active (refer to portCfgGeMediaType for more information). If the optical GbE port is active and an SFP is installed (copper or optical), "id" is displayed. If nothing is installed on the optical port, the Media field shows "--". The "cu" field also displays for interchassis links (ICLs). |
| id | Serial ID. Indicates that an SFP is installed. Use sfpShow to get more information about the SFP, including the serial number. |

Speed

The speed of the port. Valid port speeds include the following:

| | |
|-----|---|
| 1G | 1 Gbps fixed transfer speed (not supported on 16 Gbps-capable ports) |
| N1 | 1 Gbps negotiated transfer speed (not supported on 16 Gbps-capable ports) |
| 2G | 2 Gbps fixed transfer speed (only supported with use of 8 Gbps SFPs) |
| N2 | 2 Gbps negotiated transfer speed (only supported with use of 8 Gbps SFPs) |
| 4G | 4 Gbps fixed transfer speed |
| N4 | 4 Gbps negotiated transfer speed |
| 8G | 8 Gbps fixed transfer speed |
| N8 | 8 Gbps negotiated transfer speed |
| 10G | 10 Gbps fixed transfer speed |
| N10 | 10 Gbps negotiated transfer speed |
| 16G | 16 Gbps fixed transfer speed |
| N16 | 16 Gbps negotiated transfer speed |
| AN | Autonegotiating |
| UN | Unknown |

State

Port state information. Valid states include the following:

| | |
|-----------|--|
| No_Card | No interface card present. |
| No_Module | No module (GBIC or other) present. |
| Mod_Val | Module validation in process. |
| Mod_Inv | Module speed mismatch or incompatible SFP. |

| | |
|------------------|---|
| No_Light | The module is not receiving light. This state is not applicable to 16 Gbps-capable interchassis link (ICL) ports. |
| No_SigDet | No signal is detected on the port. For 16 Gbps-capable interchassis link (ICL) ports, this state replaces the No_Light indicator. It indicates that a quad small form-factor pluggable (QSFP) has been installed but is not connected with a cable. |
| No_Sync | The module is receiving light but is out of sync. |
| In_Sync | The module is receiving light and in sync. |
| Laser_Flt | The module is signaling a laser fault. |
| Port_Flt | The port is marked faulty. |
| Diag_Flt | The port failed diagnostics. |
| Lock_Ref | The port is locking to the reference signal. |
| Testing | The port is running diagnostics. |
| Offline | A port connection is not established (for virtual ports only). |
| Online | The port is up and running. |
| Transient | The port is disabled and re-enabled by MAPS to recover from the bottleneck condition caused by the target device. |
| Proto | Protocol support by GbE port. Valid protocols include the following: |
| ISCSI | The port supports ISCSI (deprecated). |
| FCIP | The port supports FCIP. |
| FCoE | The port supports Fibre Channel over Ethernet. |
| LAN | The port supports LAN. |

comment

Optionally displays one of the following:

Copper or Optical

Displays which GbE port is currently active: **Copper** indicates that the RJ45 GbE port is currently active. Only copper connections are accepted (default). **Optical** indicates that the currently active GbE port accepts both copper and optical connections (SFPs). This parameter is set by the **portCfgGeMediaType** command, and applies only to the ge0 and ge1 ports on the Brocade 7800.

Disabled

The port is disabled. Port disable reasons may be stated in parenthesis:

Disabled (FMS Mode)

The port is disabled and in Ficon Server Management mode.

Disabled (No area available for PID assignment)

The port is disabled because a PID could not be assigned for the stated reason.

Disabled (persistent) *disable reason*

This port has been disabled with the **portCfgPersistentDisable** command.

Disabled (Persistently disabled port) *disable reason*

This port has been disabled for unspecified reasons.

Disabled (Fabric Vision License required)

This D_Port has been disabled because the Fabric Vision license is not installed on the switch.

Disabled (Insistent Domain ID)

This port has been disabled because the switch did not get the requested domain ID.

Disabled (EX_Port IR POD License Limit Exceeded)

The EX_Port has been disabled because the Integrated Routing license limit for the number of EX_Ports allowed in a switch has exceeded.

Disabled (License not Installed for Integrated Routing)

The EX_port has been disabled because License for Integrated Routing not installed in the switch.

Disabled (Incompatible AMP Version)

This port has been disabled because of Analytics Monitoring Platform (AMP) version conflict.

Bypassed

The port is bypassed (loop only).

Loopback

The port is in loopback mode.

D_Port

Diagnostic port; D_Port feature can be configured on both E_Port and F_Port. D_Port on E_Port displays the world wide node name (WWNN) of the remote switch and D_Port on F_Port displays the world wide port name (WWPN) of remote port.

| | |
|---------------------------------|--|
| E_Port | Fabric port; displays the world wide name (WWN) and name of the attached switch. If the port is configured as an EX_Port, the WWN of the attached switch is the same as the router. |
| F_Port | Point-to-point port; displays the WWN of the attached N_Port. If that specific F_Port receives 1 FDISK frame, switchShow displays the total number of NPIV Public ports as 1, for example: F_Port 1 NPort + 1 NPIV devices. If the base device logs out with NPIV devices logged in, the command displays the number of NPIV devices, for example, F_Port 1 NPort + 2 NPIV. |
| G_Port | Point-to-point port, but not yet E_Port or F_Port. |
| L_Port | Loop port; displays the number of NL_Ports. |
| EX_Port | Router port; displays the WWN of the attached edge switch. |
| VF_Port | FCoE Virtual F_Port. For these ports, the number of NPIV or external device logins is displayed ("n VN-Port(s)"), instead of the WWN of the internal port. |
| VF_Port Disabled | FCoE Virtual F_Port is disabled. |
| D_Port Dynamic | The port is in Dynamic D_Port mode. Use the configure command to set the Dynamic D_Port mode. |
| D_Port On-demand | The port is in On-demand D_Port mode. Use the configure command to set the On-demand D_Port mode. |
| D_Port protocol violation | The port is stuck in G_Port state. The Static D_Port configuration is removed but the remote port is still in Static D-Port mode, which results in the switch port stuck in G_Port mode. |
| D_Port F_Port WWN Dynamic | The port is connected to a device port in the D_Port mode. |
| SIM Port | The port is a simulated (SIM) port. |
| Mirror Port | The port is a mirror port. |
| (Trunk master) | The port is the master port in a group of trunking ports. |
| (Trunk port, master is port #x) | The port is configured as a trunking port; the master port is port #x. |
| (upstream) | The E_Port is an upstream path toward the principal switch of the fabric. |

(downstream)

The E_Port is a downstream path away from the principal switch of the fabric.

FICON Persistent DID

This port has been disabled, because the switch could not obtain its configuration domain ID during the fabric reconfiguration when **fmsmode** was enabled. See the **ficonCupSet** help page for more information.

Fabric ID conflict

Two different fabrics have been assigned the same fabric ID. Applicable only to EX_Ports and Logical Fabric environments.

Fabric ID oversubscribed

One fabric has been assigned two different fabric IDs (EX_Ports only).

AoQ

Application-oriented QoS; indicates that an F_Port or N_Port has negotiated a link that is capable of quality of service (QoS). Both sides of the link have QoS capability and agreed on the protocol. The link could be between an HBA and an Access Gateway, between an Access Gateway and an edge switch, or between an HBA and an edge switch.

LB mode

If the Access Gateway cannot negotiate QoS capabilities with the edge switch, an HBA connected to the Access Gateway will not be able to negotiate a QoS link with the Access Gateway.

(logical)

Indicates a logical port. The **switchShow** output shows all logical ports currently present in the logical switch. The command displays -1 for the slot for logical ports and the user port number for slot port. The logical port numbers are not persistent and may change when the logical interswitch links (LISLs) are deleted and recreated. A logical port is shown to be in one of the following states: E_Port (if the port is online), offline, or disabled. When the port is disabled, a reason is provided.

segmented

Indicates a segmented or disabled port along with one of the following segmentation reasons:

(Encrypt incompatible)

Port segmentation or port disable due to mismatched configurations.

(Compress incompatible)

Port segmentation or port disable due to mismatched configurations.

(Encrypt limitation)

Port segmentation or port disable due to reaching encryption limitations.

(Compress limitation)

Port segmentation or port disable due to reaching compression limitations.

(Authentication failure)

Port segmentation or port disable due to authentication failure.

(Defzone conflict)

Port segmentation due to defzone states.

(Type mismatch)

D_Port configuration mismatch between local and remote switch.

(D-port mode mismatch, Not D-port)

The local port is configured as D_Port and remote port is not a D_Port. This segmentation reason is displayed for the local port.

(D-port mode mismatch, Static D-port)

The local port is configured as D_Port and remote port is not a D_Port. This segmentation reason is displayed for the remote port.

(ESC mismatch, AMP Version Conflict)

Port segmentation due to Analytics Monitoring Platform (AMP) version mismatch.

(Incompatible AMP Version)

Port segmentation due to incompatible Analytics Monitoring Platform (AMP) version.

When used with the **-slot** option, the command displays the following blade-specific information:

slot

Slot number.

Blade Type

Type of blade, for example, Core blade or AP blade. Refer to the **slotShow** command for a listing of supported blade types.

ID

A numeric blade ID that specifies the blade type. Refer to the **slotShow** command for a listing of supported blade IDs.

Status

Enabled or disabled.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

On the Brocade FCoE10-24 blade, the Proto column identifies the FCoE ports of the switch. For all FCoE ports, the speed is 10 Gbps. The default configuration of an FCoE port is an F_Port configuration. After a successful reboot **switchShow** displays all FCoE ports as online. For each FCoE port, the FCoE controller WWN is shown. FCoE ports are not configurable with the Fabric OS port commands. Use the FCoE commands instead.

If a port is configured as a long distance port, the long distance level is displayed in the format of Lx, where x represents the long distance level number. See **portCfgLongDistance** for the level description.

When a port is configured as an N_Port and is online, **switchShow** displays its type as an N_Port. Also, **switchShow** displays the WWN of the border switch attached to this N_Port as a 24-bit Port Identifier assigned to this port by the enterprise fabric.

In an AD context, if one of the L_Ports or NPIV Ports is a part of the current AD, the complete device information attached to the port is displayed.

Operands

This command has the following operands:

-slot slot

Displays blade information. You can specify this operand with **-portname** or **-qsfp**.

-portname Displays the name for each port on the switch. The port name is set by the **portName** command.

-qsfp Displays the output of the **switchShow** command with a QSFP column added. If a QSFP is installed on the port, the QSFP number is displayed for the port. This number corresponds to the physical QSFP number that can be found on the front of each Core Blade. A double dash (--) indicates that QSFP is not supported on that blade.

-portcount Displays the number of ports on the switch. This operand is exclusive.

Examples To display the port count:

```
switch:admin> switchshow -portcount
FC ports = 198, GE ports = 12
```

To display a QoS-capable Core Access Gateway with online AoQ F_Ports and N_Ports:

```
switch:admin> switchshow
switchName: Spirit_125
switchType: 66.1
switchState: Online
switchMode: Access Gateway Mode
switchWwn: 10:00:00:05:1e:85:95:d0
switchBeacon: OFF
FC Router: OFF
FC Router BB Fabric ID: 1
```

| Area | Port | Media | Speed | State | Proto |
|------|------|-------|-------|-----------|--|
| 0 | 0 | -- | N8 | No_Module | FC |
| 1 | 1 | -- | N8 | No_Module | FC |
| 2 | 2 | -- | N8 | No_Module | FC |
| 3 | 3 | -- | N8 | No_Module | FC |
| 4 | 4 | -- | N8 | No_Module | FC |
| 5 | 5 | -- | N8 | No_Module | FC |
| 6 | 6 | -- | N8 | No_Module | FC |
| 7 | 7 | -- | N8 | No_Module | FC |
| 8 | 8 | -- | N8 | No_Module | FC |
| 9 | 9 | id | N8 | Online | FC F-Port \ |
| | | | | | 10:00:00:05:1e:53:2c:54 0x690105 (AoQ) |
| 10 | 10 | id | N8 | Online | FC F-Port \ |
| | | | | | 10:00:00:05:1e:56:5f:a9 0x690107 (AoQ) |
| 11 | 11 | id | N4 | Online | FC F-Port \ |
| | | | | | 10:00:00:05:1e:56:5f:a8 0x690106 (AoQ) |
| 12 | 12 | -- | N8 | No_Module | FC |
| 13 | 13 | -- | N8 | No_Module | FC |
| 14 | 14 | -- | N8 | No_Module | FC |
| 15 | 15 | -- | N8 | No_Module | FC |
| 16 | 16 | id | N8 | No_Light | FC |
| 17 | 17 | id | N8 | Online | FC F-Port \ |
| | | | | | 1 N Port + 3 NPIV public (AoQ) |
| 18 | 18 | -- | N8 | No_Module | FC |
| 19 | 19 | -- | N8 | No_Module | FC |
| 20 | 20 | -- | N8 | No_Module | FC |
| 21 | 21 | -- | N8 | No_Module | FC |

switchShow

```

22 22  --  N8  No_Module  FC
23 23  --  N8  No_Module  FC
24 24  --  N8  No_Module  FC
25 25  --  N8  No_Module  FC
26 26  --  N8  No_Module  FC
27 27  --  N8  No_Module  FC
28 28  --  N8  No_Module  FC
29 29  --  N8  No_Module  FC
30 30  --  N8  No_Module  FC
31 31  --  N8  No_Module  FC
32 32  id  N8  Online     FC  N-Port \
    10:00:00:05:1e:43:e8:02 0x690100 (Trunk master) (AoQ)
33 33  --  N8  No_Module  FC
34 34  id  N8  Online     FC  N-Port \
    10:00:00:05:1e:43:e8:02 0x690100 (Trunk port, \
    master is Port 32 )(AoQ)
35 35  --  N8  No_Module  FC
36 36  --  N8  No_Module  FC
37 37  --  N8  No_Module  FC
38 38  --  N8  No_Module  FC
39 39  --  N8  No_Module  FC

```

To display switch information on a Virtual Fabrics-enabled switch with an assigned fabric name:

```

switch:admin> switchshow
switchName:      brocade218
switchType:      62.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    1
switchId:        fffc01
switchWwn:       10:00:00:60:69:80:04:92
zoning:          ON (testcfg1)
switchBeacon:    OFF
FC Router:       OFF
Fabric Name:     Fabric_A12
HIF Mode:ON
Allow XISL use: ON
LS Attributes:   [FID: 10, Base Switch: No,\
                 Default Switch: No, Address Mode 0]

Index Slot Port Address Media Speed State      Proto
=====
377  12  41  32f180  --  N8  No_Module
378  12  42  32f280  --  N8  No_Module
379  12  43  32f380  --  N8  No_Module
380  12  44  32f480  --  N8  No_Module
381  12  45  32f580  --  N8  No_Module
382  12  46  32f680  --  N8  No_Module
383  12  47  32f780  --  N8  No_Module
769  -1  769   --    --  --  Online  E-Port \
    10:00:00:05:1e:40:f0:79 "Switch 1" (logical)
770  -1  770   --    --  --  Offline (logical)
785  -1  785   --    --  --  Offline \
    Disabled (logical, reason why port was disabled)
(output truncated)

```


To display switch configuration information on the Brocade FCoE10-24:

```
switch:admin> switchshow -slot 8
switchName:      switch
switchType:      77.3
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:    2
switchId:        fffc02
switchWwn:       10:00:00:05:1e:75:af:00
zoning:          ON (longevity_cfg)
switchBeacon:    OFF
FC Router:       OFF
Fabric Name:     fl_2010_11_16_21_29_26__
Allow XISL Use:  OFF
LS Attributes:   [FID: 128, Base Switch: \
                  No, Default Switch: Yes, Address Mode 0]
```

| Slot | Blade Type | ID | Model Name | Status |
|------|------------|----|------------|---------|
| 8 | AP BLADE | 74 | FCOE10-24 | ENABLED |

| Index | Slot | Port | Address | Media | Speed | State | Proto |
|-------|------|------|---------|-------|-------|--------|---------------------------|
| 192 | 8 | 0 | 02c000 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 193 | 8 | 1 | 02c100 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 194 | 8 | 2 | 02c200 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 195 | 8 | 3 | 02c300 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 196 | 8 | 4 | 02c400 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 197 | 8 | 5 | 02c500 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 198 | 8 | 6 | 02c600 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 199 | 8 | 7 | 02c700 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 200 | 8 | 8 | 02c800 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 201 | 8 | 9 | 02c900 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 202 | 8 | 10 | 02ca00 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 203 | 8 | 11 | 02cb00 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 204 | 8 | 12 | 02cc00 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 205 | 8 | 13 | 02cd00 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 206 | 8 | 14 | 02ce00 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 207 | 8 | 15 | 02cf00 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 208 | 8 | 16 | 02d000 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 209 | 8 | 17 | 02d100 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 210 | 8 | 18 | 02d200 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 211 | 8 | 19 | 02d300 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 212 | 8 | 20 | 02d400 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 213 | 8 | 21 | 02d500 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 214 | 8 | 22 | 02d600 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |
| 215 | 8 | 23 | 02d700 | -- | 10G | Online | FCoE VF-Port 0 VN-Port(s) |

To display blade information and port names:

```
switch:admin> switchshow -slot 5 -portname

FC Router:       OFF
Allow XISL Use:  OFF
LS Attributes:[FID: 128, Base Switch: No, Default Switch:\
                  Yes, Address Mode 0]
```

```

Slot   Blade Type      ID      Status
-----
5      COREBLADE        52      ENABLED

Index  Slot Port      PortWWN      Name
=====
384   5     0     20:04:00:05:33:0e:df:00  MyName_portname0
385   5     1     20:04:00:05:33:0e:df:01  -----
386   5     2     20:04:00:05:33:0e:df:02  -----

```

To display QSFP information on a Brocade DCX 8510-4:

- Ports 3/60-63 indicate no QSFP is present.
- Slot 5 and 8 have QSFPs installed and the number is the group ID.
- Port 5/6 is an example of non-contiguous ports crossing QSFP group boundaries.

```

switch:admin> switchshow -qsfp
switchType:      120.1
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:    60
switchId:        fffc3c
switchWwn:       10:00:00:05:1e:40:68:78
zoning:         ON (WB_DEFAULT_CFG_LSAN)
switchBeacon:    OFF
FC Router:      OFF
Allow XISL Use: OFF
LS Attributes:  [FID: 128, Base Switch: No,
                Default Switch: Yes, Address Mode 0]

Index Slot Port QSFP  Address Media  Speed  State  Proto
=====
812   3    60   --   3c2cc0  --    N8    No_Module  FC
813   3    61   --   3c2dc0  --    N8    No_Module  FC
814   3    62   --   3c2ec0  --    N8    No_Module  FC
815   3    63   --   3c2fc0  --    N8    No_Module  FC
384   5     0   15   -----  id    16G   No_SigDet  FC
385   5     1   15   -----  id    16G   No_SigDet  FC
386   5     2    7   -----  id    16G   No_SigDet  FC
387   5     3    7   -----  id    16G   No_SigDet  FC
388   5     4    7   -----  id    16G   No_SigDet  FC
389   5     5    7   -----  id    16G   No_SigDet  FC
390   5     6   15   -----  id    16G   No_SigDet  FC
391   5     7   15   -----  id    16G   No_SigDet  FC
392   5     8   14   -----  --    16G   No_Module  FC
393   5     9   14   -----  --    16G   No_Module  FC
394   5    10    6   -----  --    16G   No_Module  FC
395   5    11    6   -----  --    16G   No_Module  FC
396   5    12    6   -----  --    16G   No_Module  FC
397   5    13    6   -----  --    16G   No_Module  FC
398   5    14   14   -----  --    16G   No_Module  FC
399   5    15   14   -----  --    16G   No_Module  FC
400   5    16   13   -----  --    16G   No_Module  FC
401   5    17   13   -----  --    16G   No_Module  FC
402   5    18    5   -----  --    16G   No_Module  FC
403   5    19    5   -----  --    16G   No_Module  FC
404   5    20    5   -----  --    16G   No_Module  FC

```

```

405 5 21 5 ----- -- 16G No_Module FC
406 5 22 13 ----- -- 16G No_Module FC
407 5 23 13 ----- -- 16G No_Module FC
408 5 24 12 ----- id 16G No_SigDet FC
409 5 25 12 ----- id 16G No_SigDet FC
410 5 26 4 ----- id 16G Online FC \
    E-Port 10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
411 5 27 4 ----- id 16G Online FC \
    E-Port 10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
412 5 28 4 ----- id 16G Online FC \
    E-Port 10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
413 5 29 4 ----- id 16G Online FC \
    E-Port 10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
414 5 30 12 ----- id 16G No_SigDet FC
415 5 31 12 ----- id 16G No_SigDet FC
1152 5 32 11 ----- -- 16G No_Module FC
1153 5 33 11 ----- -- 16G No_Module FC
1154 5 34 3 ----- -- 16G No_Module FC
1155 5 35 3 ----- -- 16G No_Module FC
1156 5 36 3 ----- -- 16G No_Module FC
1157 5 37 3 ----- -- 16G No_Module FC
1158 5 38 11 ----- -- 16G No_Module FC
1159 5 39 11 ----- -- 16G No_Module FC
1160 5 40 10 ----- -- 16G No_Module FC
1161 5 41 10 ----- -- 16G No_Module FC
1162 5 42 2 ----- -- 16G No_Module FC
1163 5 43 2 ----- -- 16G No_Module FC
1164 5 44 2 ----- -- 16G No_Module FC
1165 5 45 2 ----- -- 16G No_Module FC
1166 5 46 10 ----- -- 16G No_Module FC
1167 5 47 10 ----- -- 16G No_Module FC
1168 5 48 9 ----- -- 16G No_Module FC
1169 5 49 9 ----- -- 16G No_Module FC
1170 5 50 1 ----- -- 16G No_Module FC
1171 5 51 1 ----- -- 16G No_Module FC
1172 5 52 1 ----- -- 16G No_Module FC
1173 5 53 1 ----- -- 16G No_Module FC
1174 5 54 9 ----- -- 16G No_Module FC
1175 5 55 9 ----- -- 16G No_Module FC
1177 5 57 8 ----- id 16G No_SigDet FC
1178 5 58 0 ----- -- 16G No_Module FC
1179 5 59 0 ----- -- 16G No_Module FC
1180 5 60 0 ----- -- 16G No_Module FC
1181 5 61 0 ----- -- 16G No_Module FC
1182 5 62 8 ----- id 16G No_SigDet FC
1183 5 63 8 ----- id 16G No_SigDet FC
416 8 0 15 ----- id 16G No_SigDet FC
417 8 1 15 ----- id 16G No_SigDet FC
418 8 2 7 ----- id 16G No_SigDet FC
419 8 3 7 ----- id 16G No_SigDet FC
420 8 4 7 ----- id 16G No_SigDet FC
421 8 5 7 ----- id 16G No_SigDet FC
422 8 6 15 ----- id 16G No_SigDet FC
423 8 7 15 ----- id 16G No_SigDet FC
424 8 8 14 ----- -- 16G No_Module FC
425 8 9 14 ----- -- 16G No_Module FC
426 8 10 6 ----- -- 16G No_Module FC
427 8 11 6 ----- -- 16G No_Module FC
428 8 12 6 ----- -- 16G No_Module FC
429 8 13 6 ----- -- 16G No_Module FC

```

switchShow

```

430      8      14      14      -----      --      16G      No_Module      FC
431      8      15      14      -----      --      16G      No_Module      FC
432      8      16      13      -----      --      16G      No_Module      FC
433      8      17      13      -----      --      16G      No_Module      FC
434      8      18      5       -----      --      16G      No_Module      FC
435      8      19      5       -----      --      16G      No_Module      FC
436      8      20      5       -----      --      16G      No_Module      FC
437      8      21      5       -----      --      16G      No_Module      FC
438      8      22      13     -----      --      16G      No_Module      FC
439      8      23      13     -----      --      16G      No_Module      FC
440      8      24      12     -----      id      16G      No_SigDet      FC
441      8      25      12     -----      id      16G      No_SigDet      FC
442      8      26      4       -----      id      16G      Online         FC \
E-Port  10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
443      8      27      4       -----      id      16G      Online         FC \
E-Port  10:00:00:05:1e:39:f0:ca "ras040" (upstream)(Trunk master)
444      8      28      4       -----      id      16G      Online         FC \
E-Port  10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
445      8      29      4       -----      id      16G      Online         FC \
E-Port  10:00:00:05:1e:39:f0:ca "ras040" (Trunk master)
446      8      30      12     -----      id      16G      No_SigDet      FC
447      8      31      12     -----      id      16G      No_SigDet      FC
1184     8      32      11     -----      --      16G      No_Module      FC
1185     8      33      11     -----      --      16G      No_Module      FC
1186     8      34      3       -----      --      16G      No_Module      FC
1187     8      35      3       -----      --      16G      No_Module      FC
1188     8      36      3       -----      --      16G      No_Module      FC
1189     8      37      3       -----      --      16G      No_Module      FC
1190     8      38      11     -----      --      16G      No_Module      FC
1191     8      39      11     -----      --      16G      No_Module      FC
1192     8      40      10     -----      id      16G      No_SigDet      FC
1193     8      41      10     -----      id      16G      No_SigDet      FC
1194     8      42      2       -----      --      16G      No_Module      FC
1195     8      43      2       -----      --      16G      No_Module      FC
1196     8      44      2       -----      --      16G      No_Module      FC
1197     8      45      2       -----      --      16G      No_Module      FC
1198     8      46      10     -----      id      16G      No_SigDet      FC
1199     8      47      10     -----      id      16G      No_SigDet      FC
1200     8      48      9       -----      --      16G      No_Module      FC
1201     8      49      9       -----      --      16G      No_Module      FC
1202     8      50      1       -----      --      16G      No_Module      FC
1203     8      51      1       -----      --      16G      No_Module      FC
1204     8      52      1       -----      --      16G      No_Module      FC
1205     8      53      1       -----      --      16G      No_Module      FC
1206     8      54      9       -----      --      16G      No_Module      FC
1207     8      55      9       -----      --      16G      No_Module      FC
1208     8      56      8       -----      id      16G      No_SigDet      FC
1209     8      57      8       -----      id      16G      No_SigDet      FC
1210     8      58      0       -----      --      16G      No_Module      FC
1211     8      59      0       -----      --      16G      No_Module      FC
1212     8      60      0       -----      --      16G      No_Module      FC
1213     8      61      0       -----      --      16G      No_Module      FC
1214     8      62      8       -----      id      16G      No_SigDet      FC
1215     8      63      8       -----      id      16G      No_SigDet      FC

```

To display media type information (relevant output excerpts only):

- The following example shows **switchshow** output for the ge0 and ge1 ports on a Brocade 7800. The ge0 port is a copper port (default). The ge1 has an SFP installed (copper or optical):

```
[...]
ge0 cu    1G  No_Sync  FCIP  Copper
ge1 id    1G  No_Light FCIP
ge2 --    1G  No_Module FCIP
ge3 --    1G  No_Module FCIP
ge4 id    1G  Online   FCIP
ge5 --    1G  No_Module FCIP
[...]
```

- This example shows **switchshow** output for a Brocade 5100. Port 31 has an SFP installed. Nothing is installed on Port 32.

```
[...]
31 31  421f00  id    N4  In_Sync  FC  Disabled
32 32  422000  --    N8  No_Module FC  Disabled
[...]
```

- The **sfpShow** output for the same switch displays the serial number for the SFP.

```
[...]
Port 31: id (sw) Vendor: BROCADE Serial No: UAF1081800000MK
        Speed: 200,400,800_MB/s
Port 32: --
[...]
```

To display output when an authentication fails for a reason and the port gets disabled:

```
[...]
15 15  id    N2  No_Light          Disabled (Authentication Required)
[...]
```

To display the neighbor switch WWN for the segmented ISLs during ELP and post ELP:

```
switch:admin> switchshow
switchName:    swd77
switchType:    71.2
switchState:   Online
switchMode:    Native
switchRole:    Principal
switchDomain:   1
switchId:      fffc01
switchWwn:     10:00:00:05:1e:a2:ec:9c
zoning:        OFF
switchBeacon:  OFF
```

```
Index Port Address Media Speed State      Proto
=====
  0  0  010000  id    N8  No_Light  FC  Disabled (Persistent)
  1  1  010100  id    N8  No_Light  FC  Disabled (Persistent)
  2  2  010200  id    N4  No_Light  FC
  3  3  010300  id    N4  No_Light  FC
  4  4  010400  id    N8  No_Light  FC  Disabled (Persistent)
  5  5  010500  id    N8  No_Light  FC  Disabled (Persistent)
  6  6  010600  id    N8  No_Light  FC  Disabled (Persistent)
```

switchShow

```

 7  7  010700  id    N8  No_Light  FC  Disabled (Persistent)
 8  8  010800  --   N8  No_Module FC  Disabled (Persistent)
 9  9  010900  id    N4  Online   FC  E-Port \
      10:00:00:05:1e:a3:00:59 segmented,(RA TOV incompat)
10 10  010a00  id    N8  No_Light  FC  Disabled (Persistent)

```

To display the output when any of the ICL links get disabled with no Enterprise ICL (EICL) license installed:

- If no EICL license is installed and the number of ICL-connected chassis in the fabric is more than 4, the port disable reason code displayed is, "No EICL License".
- If EICL license is installed and the number of ICL-connected chassis in the fabric exceeds 10, the port disable reason code displayed is, "EICL License Limited".

```

switch:admin> switchshow
switchName:      DCX_150
switchType:      120.3
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     6
switchId:        fffc06
switchWwn:       10:00:00:05:1e:93:ff:00
zoning:          OFF
switchBeacon:    OFF
FC Router:       OFF
Fabric Name:     fabric
Allow XISL Use:  OFF
LS Attributes:   [FID: 128, Base Switch: No, Default Switch: Yes, Address Mode 0]

```

| Index | Slot | Port | Address | Media | Speed | State | Proto | |
|-----------------|------|------|---------|-------|-------|-----------|-------|--|
| 0 | 1 | 0 | 050000 | -- | N8 | No_Module | FC | |
| 1 | 1 | 1 | 050100 | -- | N8 | No_Module | FC | |
| 2 | 1 | 2 | 050200 | -- | N8 | No_Module | FC | |
| | | | | | | | | |
| 398 | 5 | 14 | ----- | cu | 8G | No_Sync | FC | |
| 399 | 5 | 15 | ----- | cu | 8G | No_Sync | FC | |
| 400 | 5 | 16 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 401 | 5 | 17 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 402 | 5 | 18 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 403 | 5 | 19 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 404 | 5 | 20 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 405 | 5 | 21 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 406 | 5 | 22 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| 407 | 5 | 23 | ----- | cu | 8G | In_Sync | FC | Disabled (No EICL License) |
| | | | | | | | | |
| 416 | 8 | 0 | ----- | cu | 8G | In_Sync | FC | E-Port 10:00:00:05:1e:4a:cb:00 "sw0" (upstream) \ |
| (Trunk master) | | | | | | | | |
| 417 | 8 | 1 | ----- | cu | 8G | In_Sync | FC | E-Port (Trunk port, master is |
| Slot 8 Port 0) | | | | | | | | |
| 418 | 8 | 2 | ----- | cu | 8G | In_Sync | FC | E-Port (Trunk port, master is |
| Slot 8 Port 0) | | | | | | | | |
| 419 | 8 | 3 | ----- | cu | 8G | In_Sync | FC | E-Port (Trunk port, master is |
| Slot 8 Port 0) | | | | | | | | |
| 420 | 8 | 4 | ----- | cu | 8G | In_Sync | FC | E-Port (Trunk port, master is |
| Slot 8 Port 0) | | | | | | | | |

To display the port status information on 16 Gbps-capable Inter-Chassis Links (ICLs):

```
switch:admin> switchshow
switchName:      ICL_DCX
switchType:      120.3
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:05:1e:48:f8:02
zoning:          ON (bb_zone)
switchBeacon:    OFF
FC Router:       ON
Fabric Name:     BBTI
LS Attributes:   [FID: 10, Base Switch: Yes,
Default Switch: No, Address Mode 0]

Index Slot Port Address Media Speed State Proto
=====
396 5 12 018040 id 16G Online FC EX-Port \
10:00:00:05:33:b3:39:00 "Pluto2" (fabric id = 11 )(Trunk master)
397 5 13 018080 id 16G Online FC EX-Port \
10:00:00:05:33:b3:39:00 "Pluto2" (fabric id = 11 )(Trunk master)
398 5 14 0180c0 id 16G Online FC EX-Port \
10:00:00:05:33:b3:39:00 "Pluto2" (fabric id = 11 )(Trunk master)
399 5 15 018100 id 16G Online FC EX-Port \
10:00:00:05:33:b3:39:00 "Pluto2" (fabric id = 11 )(Trunk master)
412 5 28 ----- id 16G Online FC E-Port \
10:00:00:05:1e:75:4c:00 "FCR-DCX128" (upstream)(Trunk master)
413 5 29 ----- id 16G Online FC E-Port \
10:00:00:05:1e:75:4c:00 "FCR-DCX128" (Trunk master)
414 5 30 ----- id 16G Online FC E-Port \
10:00:00:05:1e:75:4c:00 "FCR-DCX128" (Trunk master)
415 5 31 ----- id 16G Online FC E-Port \
10:00:00:05:1e:75:4c:00 "FCR-DCX128" (Trunk master)
```

To display the output when the location ID is configured:

```
switch:admin> switchshow
switchName:      switch
switchType:      66.1
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     106
switchId:        fffc6a
switchWwn:       10:00:00:05:1e:57:df:49
zoning:          ON (cfg_qos)
switchBeacon:    OFF
FC Router:       ON
FC Router BB Fabric ID: 128
Address Mode:    0
```

switchShow

```
Fabric Name:    base
Location ID:    1

Index Port Address Media Speed State      Proto
=====
```

SEE ALSO [fabricName](#), [portCfgLongDistance](#), [switchDisable](#), [switchEnable](#), [switchName](#)

switchUptime

Displays the amount of time the switch has been operating.

Synopsis **switchuptime**

Description Use this command to display the current time and the amount of time that the switch has been operational.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To view the uptime for the switch:

```
switch:user> switchuptime  
9:50pm   up for 20 mins
```

SEE ALSO None

switchViolation

Dumps the DCC violations for a switch.

Synopsis `switchViolation --dump -dcc`

Description Use this command to display all Device Connection Control (DCC) violations that have occurred on a switch. Internally the command searches "errdumpall" for the DCC violations. For each DCC violation, the command displays the device WWN and the port where the violation occurred.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command can be executed on both active and standby CPs. This command does not support High Availability (HA).

Operands This command has the following operands. If executed without operands, the command prints the usage.

--dump Displays specified policy violation.

-dcc Specifies the violation type as DCC.

Examples To display DCC violations for a switch:

```
switch:admin> switchviolation --dump -dcc
Device WWN                               Port
-----
22:00:00:04:cf:75:59:87                   10
```

SEE ALSO None

syslogAdmin

Configures a syslog server host.

Synopsis `syslogadmin --set -ip ip_address | hostname [-secure [-port port_num]]`

`syslogadmin --set -facility level`

`syslogadmin --remove -ip ip_address | hostname`

`syslogadmin --show -ip | -facility`

`syslogadmin --help`

Description Use this command to configure a switch to forward all error log entries to a remote syslog server, to set the syslog facility to a specified log file, to remove a syslog server, and to display the list of configured syslog servers. Brocade switches use the syslog daemon, a process available on most UNIX systems that reads and forwards system messages to the appropriate log files or users, depending on the system configuration. Up to six servers are supported.

By default, the switch uses UDP protocol to send the error log messages to the syslog server. The default UDP port is 514. Use the **-secure** option to configure the switch to send the error log messages securely using the Transport Layer Security (TLS) protocol. TLS is an encryption protocol over the TCP/IP network protocol and it can be used only with the TCP-based destinations (`tcp()` and `tcp6()`). The default TLS port is 6514. While enabling secure syslog mode, you must specify a port that is configured to receive the log messages from the switch.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--set -ip *ip_address* | *hostname*

Configures a syslog server with the specified IP address or hostname. IPv4 and IPv6 addresses are supported. Only one syslog server can be specified at any given time. To configure more than one server, the command must be executed for each server.

-secure [-port *port_num*]

Enables secure syslog mode to send the error log messages securely using the TLS protocol to the syslog server. The secure syslog mode is disabled by default. The **-port** operand is optional. If secure mode is enabled and port number is not specified, the default TLS port number (6514) is set.

--set -facility *level*

Sets the syslog facility. Valid levels are 0 through 7. The default is 7.

0

LOG_LOCAL0

1

LOG_LOCAL1

| | |
|---|----------------------|
| 2 | LOG_LOCAL2 |
| 3 | LOG_LOCAL3 |
| 4 | LOG_LOCAL4 |
| 5 | LOG_LOCAL5 |
| 6 | LOG_LOCAL6 |
| 7 | LOG_LOCAL7 (default) |

- remove -ip *ip_address* | *hostname***
Removes the specified syslog server. IPv4 and IPv6 addresses are supported.
- show**
Displays the list of configured syslog servers and the facility level.
- ip**
Displays all syslog server IP addresses and hostnames.
- facility**
Display the configured syslog facility.
- help**
Displays the command usage.

Examples

To configure an IPv4 secure syslog server to which error log messages are sent:

```
switch:admin> syslogadmin --set -ip 172.26.26.173 -secure -port 2000
```

To configure an IPv6 non-secure syslog server:

```
switch:admin> syslogadmin --set -ip fec0:60:69bc:92:218:8bff:fe40:15c4
```

To configure a syslog server using a hostname:

```
switch:admin> syslogadmin --set -ip win2k8-58-113
```

To set the syslog facility to LOG_LOCAL2:

```
switch:admin> syslogadmin --set -facility 2
switch:admin> syslogadmin --show -facility
Syslog facility: LOG_LOCAL2
```

To display all syslog IP addresses configured on a switch:

```
switch:admin> syslogadmin --show -ip
syslog.1 172.26.26.173
syslog.2 fec0:60:69bc:92:218:8bff:fe40:15c4
syslog.3 win2k8-58-113
```

To remove the IP address fec0:60:69bc:92:218:8bff:fe40:15c4 from the list of servers to which error log messages are sent:

```
switch:admin> syslogadmin --remove -ip fec0:60:69bc:92:218:8bff:fe40:15c4
```

SEE ALSO None

sysShutDown

Provides a graceful shutdown to protect the switch file systems.

Synopsis `sysshutdown`

Description On standalone platforms, use this command to shut down the switch operating system.

On enterprise-class platforms, when **sysShutDown** is called on the active control processor (CP), the command shuts down the active CP, standby CP, and any AP blades.

Some platforms will only shut down the operating systems; others will shut down the operating system as well as shut off the power, that is, the LEDs will turn black.

After executing this command, manually power off the system. To reboot the system, manually turn the power switch on.

Notes This command is not supported on the standby CP.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To perform a system shutdown on a standalone platform:

```
switch:admin> sysshutdown
This command will shutdown the operating systems on your
switch. You are required to power-cycle the switch in
order to restore operation.
Are you sure you want to shutdown the switch [y/n]? y
Broadcast message from root (ttyS0) Mon Sep 12 17: \
52:12 2005...

The system is going down for system halt NOW !!
INIT: Switching to runlevel:
INIT: Sending processes the TERM signal
ess095:root> Unmounting all filesystems.
The system is halted
flushing ide devices: hda
Power down.
```

To perform a system shutdown on a Brocade DCX-4S:

```
switch:admin> sysshutdown
This command will shutdown the operating systems on your
switch. You are required to power-cycle the switch in
order to restore operation.
Are you sure you want to shutdown the switch [y/n]?y
HA is disabled
Shutting down blade in slot:1, IP addr:127.1.14.2
Shutting down blade in slot:8, IP addr:127.1.14.9
```

```
Shutting down OCP at:0.0.0.0
```

```
Broadcast message from root (pts/0) Wed Nov 5 19:03:06 2008...
```

```
The system is going down for system halt NOW !!
```

To attempt a system shutdown from the standby CP (not supported):

```
switch:admin> sysshutdown
```

```
Shut down the whole system is not support from the standby CP
```

```
For shut down the whole system
```

```
please run the sysshutdown from the active CP
```

SEE ALSO [haDisable](#)

tempShow

Displays temperature readings.

Synopsis `tempshow [-details]`

Description Use this command to display the current temperature readings of all temperature sensors in a switch. For each sensor, this command displays the sensor ID (an index number), the slot number (if applicable), the sensor index (when issued with the `--detail` option), the sensor state (OK or absent), and the temperature. The temperature readings are given in both Centigrade and Fahrenheit.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Refer to the hardware reference manual for your switch to determine the normal temperature range.

Operands This command has the following operand:

-details Displays an additional column for the sensor index. Sensor index indicates the position of the sensor in the system.

Examples To display temperature and status sensors:

```
switch:user> tempshow
Sensor ID      Slot      State      Centigrade      Fahrenheit
=====
1              1         Ok         38              100
2              1         Ok         28              82
3              1         Ok         40              104
4              1         Ok         31              87
5              1         Ok         43              109
6              2         Ok         39              102
7              2         Ok         28              82
8              2         Ok         40              104
9              2         Ok         30              86
10             2         Ok         43              109
```

To display temperature and status sensors with the sensor index:

```
switch:user> tempshow -details
Sensor ID|Slot|Sensor Index|State|Centigrade|Fahrenheit|
-----|-----|-----|-----|-----|-----|
1|1|0|Ok|35|95|
2|1|1|Ok|34|93|
3|1|2|Ok|43|109|
4|1|3|Ok|34|93|
5|1|4|Ok|41|105|
```


| | | | | | | |
|----|---|---|--------|----|-----|--|
| 6 | 2 | 0 | Absent | 0 | 0 | |
| 7 | 4 | 0 | Ok | 33 | 91 | |
| 8 | 4 | 1 | Ok | 42 | 107 | |
| 9 | 4 | 2 | Ok | 42 | 107 | |
| 10 | 4 | 3 | Ok | 44 | 111 | |

SEE ALSO [fanShow](#), [psShow](#), [sensorShow](#), [slotShow](#)

timeOut

timeOut

Sets or displays the idle timeout value for a login session.

Synopsis `timeout [timeval]`

Description Use this command without an operand to display the current timeout value (in minutes) after which idle logins are automatically terminated.

Use this command with the *timeval* operand to set the login timeout value to the specified interval. A value of 0 disables timeout of login sessions.

The new timeout value takes effect with the next logins.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

timeval

Specify the number of minutes for the Telnet timeout value. Valid values are 1 to 99,999, or 0 to disable login timeouts. This operand is optional; if omitted, the command displays the current timeout value.

Examples To set the idle timeout to 10 minutes:

```
switch:admin> timeout 10
      IDLE Timeout Changed to 10 minutes
The modified IDLE Timeout will be in effect after NEXT login
```

SEE ALSO None

topologyShow

Displays the unicast fabric topology.

Synopsis `topologyshow -index [domain]`
`topologyshow -nopage [domain]`
`topologyshow --help`

Description Use this command to display the fabric topology as it appears to the local switch. The display varies depending on the hardware configuration. The following rules apply:

- On all switches, the command displays the number of domains in the fabric and the local Domain IDs. If translate domains are configured, existing translate domains and associated ports are displayed.
- On an edge fabric, the command displays the following additional details for all domains in the fabric (including local translate domains):
 - All possible paths from the local switch to each of the remote domains.
 - For each path, the cost, the associated output port on the local switch, the path cost, and the number of hops from the local switch to the destination switch.
 - A summary of all ports that are routed through that path.
- On a backbone fabric, the command displays details for remote domains only. Details for local translate domains are not displayed.
- If there are two switches in the Backbone and the edge fabric is directly connected to both of those switches, **topologyshow** does not display the description of the translate domain associated with that edge fabric. In this case the translate domain is considered local to both of the switches in the backbone.
- If there is only one switch in the backbone, no domain details are displayed (all domains are local).

Depending on the fabric, the display may contains the following fields:

Local Domain ID

The domain number of local switch.

Local Translate Domain x owned by port

The port number associated with the local translate domain x.

Domain

The domain number of destination switch.

Metric

The cost of reaching destination domain.

Name

The name of the destination switch.

Path Count

The number of currently active paths to the destination domain. The maximum number of paths supported is 16. If there are more than 16 paths, the path count displays as 16 of *number of paths available*, for example, "16 of 20".

Hops

The maximum number of hops to reach destination domain.

Out Port

The port to which incoming frames are forwarded to reach the destination domain.

In Ports

The input ports that use the corresponding out port to reach the destination domain. This is the same information provided by **portRouteShow** and **uRouteShow** but in a different format.

Total Bandwidth

The maximum bandwidth of the out port. A bandwidth that is less than 0.512 Gbps is adjusted to the nearest power of 2 value. A bandwidth in the range of 0.512 Gbps (included) to 1 Gbps (not included) is adjusted to the 0.512 Gbps value. No adjustment takes place if the bandwidth is greater or equal to 1 Gbps.

Bandwidth Demand

The maximum bandwidth demand by the in ports.

Flags

Always D, indicating a dynamic path. A dynamic path is discovered automatically by the FSPF path selection protocol.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

The following operand is optional:

domain

Specify the destination domain for which to display the topology information.

index

Display port index numbers instead of slot or port numbers in the output.

nopage

Disables the paging format to display the output.

Examples

To display the topology on a single switch: domain is local, details are not displayed.

```
switch:user> topologyshow

1 domains in the fabric; Local Domain ID: 97
```

When executed from an edge fabric, **topologys**how displays details for all domains, including local domains and local translate domains:

```
switch:user> topologyshow

6 domains in the fabric; Local Domain ID: 7

Domain:          1
Metric:          10500
Name:            fcr_xd_1_1
Path Count:      1

Hops:            2
Out Port:        11
In Ports:        0 1 2 3 4 5 6 7 8 9 15
```

```

    Total Bandwidth:      8.000 Gbps
    Bandwidth Demand:    1275 %
    Flags:                D

Domain:      2
Metric:     1000
Name:      fcr_fd_2
Path Count: 1

    Hops:      2
    Out Port:  15
    In Ports:  0 1 2 3 4 5 6 7 8 9 11
    Total Bandwidth: 2.000 Gbps
    Bandwidth Demand: 4000 %
    Flags:      D

Domain:      3
Metric:     10500
Name:      fcr_xd_3_5
Path Count: 1

    Hops:      2
    Out Port:  11
    In Ports:  0 1 2 3 4 5 6 7 8 9 15
    Total Bandwidth: 8.000 Gbps
    Bandwidth Demand: 1275 %
    Flags:      D

Domain:      111
Metric:     500
Name:      peng3900101
Path Count: 1

    Hops:      1
    Out Port:  15
    In Ports:  0 1 2 3 4 5 6 7 8 9 11
    Total Bandwidth: 2.000 Gbps
    Bandwidth Demand: 4000 %
    Flags:      D
(output truncate)

```

The command is executed from the backbone in a fabric with five switches. The fabric has five domains, but details are only shown for the three remote domains, not for the two local translate domains.

```
switch:user> topologyshow
```

```

5 domains in the fabric; Local Domain ID: 2
Local Translate Domain 4 owned by port: 24
Local Translate Domain 5 owned by port: 23 33

Domain:      1
Metric:     500
Name:      pengsaturn104
Path Count: 1

    Hops:      1
    Out Port:  0
    In Ports:  23 24 33 38 39
    Total Bandwidth: 8.000 Gbps
    Bandwidth Demand: 350 %

```

topologyShow

```
Flags: D
Domain: 3
Metric: 10500
Name: fcr_xd_3_6
Path Count: 1

Hops: 2
Out Port: 0
In Ports: 23 24 33 38 39
Total Bandwidth: 8.000 Gbps
Bandwidth Demand: 350 %
Flags: D
```

To display the topology with index numbers:

```
switch:user> topologyshow -index
1 domain(s) in the fabric; Local Domain ID: 188

Domain: 178
Metric: 500
Name: dcx_178
Path Count: 1

Hops: 1
Out Port (index): 112
In Ports (index): 246 247
Total Bandwidth: 24.000 Gbps
Bandwidth Demand: 133 %
Flags: D
```

SEE ALSO [fcrXlateConfig](#), [portRouteShow](#), [uRouteShow](#)

traceDump

Initiates, or removes a trace dump or displays the trace dump status.

Synopsis `tracedump [-S]`

`tracedump -n [-s s/of]`

`tracedump -r [-s s/of] | -R`

`tracedump -c`

Description

Use this command to initiate a background trace dump, to remove the content of a trace dump, or to display the dump status on the switch.

When executed without operands, this command defaults to **traceDump -S**.

Execution of **traceDump -n** generates a local trace dump locally. Use **supportSave** to transfer the local trace dump to a remote host. When **supportSave** is used, the default remote file name format for the trace dump file is as follows:

`chassisname-S#xxs-YYYYMMDDHHMMSS.type_[NEW | OLD].dmp.gz`

When the **traceDump -n -s** command is used with an AP blade, the trace dump is generated with the following file name format:

`chassisname-S#xxs-YYYYMMDDHHMMSS_dump.tar.gz`

S#

Indicates the slot number (0 on standalone platforms)

xxs

Indicates the processor. Values are cp0, cp1, bp0, bp1, dp0, or dp1.

YYYYMMDDHHMMSS

Indicates the trace dump time stamp (year-month-day-hour-minute-second).

type

Indicates the tracedump type. Valid types include TRACE, BTRACE, DTRACE, CTRACE, and Reboot.

NEW | OLD

Indicates a new or an old trace dump file.

dmp.gz

The compressed trace dump archive file extension.

For example, the filename for a new BFOS tracedump received from slot 10 on chassisABC would be `chassisABC-S10bp-201103031111.BTRACE_NEW.dmp.gz`.

traceDump

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following mutually exclusive operands:

- S** Displays the trace dump status This operand is optional. If omitted, the same status information is displayed.
- n** Initiates a background trace dump.
- r** Clears the status of a particular trace dump on a specified slot.
- R** Clears the status of a particular trace dump on all slots.
- c** Clears all trace dump buffers. This operation resets the trace buffer contents.
- s *slot*** Specifies the slot number from which a trace dump is generated. If a slot is not supplied, the trace dump is generated from the local slot. This operand is optional.

Examples To initiate a background trace dump from slot 5:

```
switch:admin> tracedump -n -s 5
```

To display the trace dump status on the Brocade 8510-8 switch:

```
switch:admin> tracedump -S
Dump status for switch:
Slot 6: Thu Mar 17 04:43:44 2016
Slot 7: Thu Mar 17 04:44:57 2016
Slot 11: Thu Mar 17 04:45:07 2016
```

To display the trace dump status on a Gen6 switch:

```
switch:admin> tracedump -S
Dump status for switch:

slot: 1
Type          Timestamp
-----
Reboot        2016/03/16 10:50
FFDC(EM-1100) 2016/03/17 02:13
Panicdump     2016/03/17 02:24
CLI           2016/03/17 02:47

slot: 2
Type          Timestamp
-----
Reboot        2016/03/16 10:50
Reboot        2016/03/17 02:24
Reboot        2016/03/17 02:47
```


To clear the status of a trace dump:

```
switch:admin> tracedump -R  
trace dump removed
```

To clear the content of the trace dump buffer:

```
switch:admin> tracedump -c  
Cleared Trace Buffer contents
```

SEE ALSO [supportFtp](#), [supportSave](#), [supportShow](#)

trunkDebug

Debugs a trunk link failure.

Synopsis `trunkdebug port1 port2`

Description Use this command to debug a trunk link failure. This command reports one of the following messages, based on the trunking properties of the two specified ports:

- Switch does not support trunking
- Trunking license required
- port *port_id* is not E_Port
- port *port_id* trunking disabled
- port *port_id* speed is not 2G, 4G or 8G
- port *port_id* and port *port_id* are not in same port group
- port *port_id* and port *port_id* connect to different switches
- port *port_id* and port *port_id* connect to the switch WWN
- port *port_id* is not a trunking port due to: E_Port being disabled, or trunking might be disabled at remote port
- port *port_id* and port *port_id* cannot trunk, please check link length to make sure difference is less than 400 m

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

port1

Specify the port index number of port 1. Use the **switchShow** command to view the index numbers for a port. This operand is required.

port2

Specify the port index number of port 2. Use the **switchShow** command to view the index numbers for a port. This operand is required.

Examples To debug a trunk connection:

```
switch:admin> trunkdebug 43 44
Switch does not support trunking
```

```
switch:admin> trunkdebug 62 63
port 62 and 63 are trunked together
```

SEE ALSO [portCfgTrunkPort](#), [switchCfgTrunk](#), [trunkShow](#)

trunkShow

Displays trunking information.

Synopsis `trunkshow [-perf] [-swname]`

Description Use this command to display trunking information of both E_Ports and EX_Ports. The command displays the following fields:

Trunking Group Number

Displays each trunking group on a switch. All ports that are part of this trunking group are displayed.

Port to port connections

Displays the port-to-port trunking connections.

WWN

Displays the world wide name of the connected port.

Domain

Displays the domain IDs of the switches directly connected to the physical ports. In case of an FC Router backbone fabric interlinking several edge fabrics, the domain ID displayed for an E_Port trunk refers to a domain of a switch within the backbone fabric, whereas the domain ID displayed for an EX_Port trunk refers to the domain ID of a switch in the edge fabric. Because they are independent fabrics, it is possible that both the backbone and the edge fabric may have the same domain ID assigned to switches. If this is the case, run **switchShow** to obtain information on the port types of the local switch and the WWNs of the remote switches. Refer to the Example section for an illustration.

deskew

The difference between the time it takes for traffic to travel over each ISL compared to the time it takes through the shortest ISL in the group plus the minimum deskew value. The value is expressed in nanoseconds divided by 10. The firmware automatically sets the minimum deskew value for the shortest ISL, which is 15.

Master

Displays whether this trunking port connection is the master port connection for the trunking group.

When used with the **-perf** option, the command output displays the following additional information:

bandwidth

The bandwidth (Rx, Tx, and the combined total for Tx+Rx) of the trunk group. Values are displayed as either bits per second (Bps), kilobits per second (Kbps), megabits per second (Mbps), or gigabits per second (Gbps), rounded down to the next integer.

throughput

Displays the throughput (Rx, Tx, and the combined total for Tx+Rx) of the trunk group. Results are displayed for the previous second. Values are displayed as either bits per second (Bps), kilobits per second (Kbps), megabits per second (Mbps), or gigabits per second (Gbps), rounded down to the next integer.

%

Displays the percentage of link utilization (Rx, Tx, and the combined total for Tx+Rx) . Even when the link utilization is 100% , the throughput value will be lesser than the bandwidth value, due to the 8b/10b encoding and the control words transmitted. For example, the throughput for an 8Gbps link at 100% utilization would be approximately 6.8Gbps.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

-perf

Displays the total bandwidth, throughput, and percentage of link utilization information for the trunk group (Rx, Tx, and combined total for Tx+Rx). This operand is optional.

-swname

Displays the neighbor switch name. This operand is optional.

Examples To display trunking information for a switch:

```
switch:admin> trunkshow
1: 43-> 0 10:00:00:05:1e:53:e3:8a 92 deskew 15 MASTER

2: 58-> 66 10:00:00:05:1e:4f:eb:00 65 deskew 16 MASTER
61-> 69 10:00:00:05:1e:4f:eb:00 65 deskew 16
57-> 65 10:00:00:05:1e:4f:eb:00 65 deskew 16
60-> 68 10:00:00:05:1e:4f:eb:00 65 deskew 15
56-> 64 10:00:00:05:1e:4f:eb:00 65 deskew 16
63-> 71 10:00:00:05:1e:4f:eb:00 65 deskew 16
62-> 70 10:00:00:05:1e:4f:eb:00 65 deskew 16

3: 59-> 67 10:00:00:05:1e:4f:eb:00 65 deskew 15 MASTER
```

To display trunking information for a switch that is part of an FC Router backbone fabric interlinking several edge fabrics (see the EX_Port with WWN "10:00:00:05:1e:35:b3:03" and the E_Port with WWN "10:00:00:05:1e:37:12:13" in the output below):

```
switch:admin> trunkshow
4: 49-> 0 10:00:00:05:1e:35:b3:03 4 deskew 16 MASTER
54-> 2 10:00:00:05:1e:35:b3:03 4 deskew 16
53-> 5 10:00:00:05:1e:35:b3:03 4 deskew 16
50-> 6 10:00:00:05:1e:35:b3:03 4 deskew 15
51-> 4 10:00:00:05:1e:35:b3:03 4 deskew 16
52-> 7 10:00:00:05:1e:35:b3:03 4 deskew 67
55-> 3 10:00:00:05:1e:35:b3:03 4 deskew 16
48-> 1 10:00:00:05:1e:35:b3:03 4 deskew 15

5: 71-> 22 10:00:00:05:1e:37:12:13 4 deskew 17 MASTER
67-> 17 10:00:00:05:1e:37:12:13 4 deskew 16
70-> 20 10:00:00:05:1e:37:12:13 4 deskew 16
69-> 21 10:00:00:05:1e:37:12:13 4 deskew 16
66-> 18 10:00:00:05:1e:37:12:13 4 deskew 17
68-> 23 10:00:00:05:1e:37:12:13 4 deskew 17
64-> 16 10:00:00:05:1e:37:12:13 4 deskew 15
65-> 19 10:00:00:05:1e:37:12:13 4 deskew 16
```

```

switch:admin> switchshow
 48 4 0 013000 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 49 4 1 013100 id N4 Online EX-Port \
    10:00:00:05:1e:35:b3:03 "SW4100_33" (fabric id = 100 )\
    (Trunk master) E-Port \
    50:00:51:e3:70:bb:af:c1 "fcr_xd_9_100"
 50 4 2 013200 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 51 4 3 013300 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 52 4 4 013400 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 53 4 5 013500 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 54 4 6 013600 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 55 4 7 013700 id N4 Online EX-Port \
    (Trunk port, master is Slot 4 Port 1 )
 64 7 0 014000 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 65 7 1 014100 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 66 7 2 014200 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 67 7 3 014300 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 68 7 4 014400 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 69 7 5 014500 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 70 7 6 014600 id N4 Online E-Port \
    (Trunk port, master is Slot 7 Port 7 )
 71 7 7 014700 id N4 Online E-Port \
    10:00:00:05:1e:37:12:13 "SW4900_43" (downstream)(Trunk master)

```

To display trunking information along with bandwidth throughput information:

```

switch:admin> trunkshow -perf
1: 43-> 0 10:00:00:05:1e:53:e3:8a 92 deskew 15 MASTER
Tx: Bandwidth 4.00Gbps, Throughput 288.00bps (0.00%)
Rx: Bandwidth 4.00Gbps, Throughput 320.00bps (0.00%)
Tx+Rx: Bandwidth 8.00Gbps, Throughput 608.00bps (0.00%)

2: 58-> 66 10:00:00:05:1e:4f:eb:00 65 deskew 16 MASTER
61-> 69 10:00:00:05:1e:4f:eb:00 65 deskew 16
57-> 65 10:00:00:05:1e:4f:eb:00 65 deskew 16
60-> 68 10:00:00:05:1e:4f:eb:00 65 deskew 15
56-> 64 10:00:00:05:1e:4f:eb:00 65 deskew 16
63-> 71 10:00:00:05:1e:4f:eb:00 65 deskew 16
62-> 70 10:00:00:05:1e:4f:eb:00 65 deskew 16
Tx: Bandwidth 28.00Gbps, Throughput 320.00bps (0.00%)
Rx: Bandwidth 28.00Gbps, Throughput 1.73Kbps (0.00%)
Tx+Rx: Bandwidth 56.00Gbps, Throughput 2.05Kbps (0.00%)

3: 59-> 67 10:00:00:05:1e:4f:eb:00 65 deskew 15 MASTER
Tx: Bandwidth 8.00Gbps, Throughput 0.00bps (0.00%)
Rx: Bandwidth 8.00Gbps, Throughput 0.00bps (0.00%)
Tx+Rx: Bandwidth 16.00Gbps, Throughput 0.00bps (0.00%)

```

trunkShow

To display trunking information along with switch name:

```
switch:admin> trunkshow -swname
1: 0-> 10 10:00:00:05:1e:a1:99:09 88 SW_88 deskew 15 MASTER

2: 2-> 11 10:00:00:05:1e:a1:eb:39 75 switchname01234567890123456789 deskew 15
  MASTER
```

To display trunking information, with switch name and bandwidth throughput information:

```
switch:admin> trunkshow -pref -swname
1: 0-> 10 10:00:00:05:1e:a1:99:09 88 SW_88 deskew 15 MASTER
  Tx: Bandwidth 4.00Gbps, Throughput 0.00bps (0.00%)
  Rx: Bandwidth 4.00Gbps, Throughput 0.00bps (0.00%)
  Tx+Rx: Bandwidth 8.00Gbps, Throughput 0.00bps (0.00%)

2: 2-> 11 10:00:00:05:1e:a1:eb:39 75 switchname01234567890123456789 deskew 15
  MASTER
  Tx: Bandwidth 4.00Gbps, Throughput 0.00bps (0.00%)
  Rx: Bandwidth 4.00Gbps, Throughput 0.00bps (0.00%)
  Tx+Rx: Bandwidth 8.00Gbps, Throughput 0.00bps (0.00%)
```

SEE ALSO [portCfgTrunkPort](#), [switchCfgTrunk](#)

tsClockServer

Displays or sets the Network Time Protocol (NTP) Server addresses.

Synopsis `tsclockserver [ipaddr [; ipaddr ...]]`

Description Use this command to synchronize the local time of the Principal or Primary FCS switch to one or more external NTP servers.

This command accepts a list of NTP server addresses. The NTP server addresses can be passed in either IPV4 or IPV6 address format or as a DNS server name. When multiple NTP server addresses are specified, **tsClockServer** sets the first reachable address for the active NTP server. The remaining addresses are stored as backup servers, which can take over if the active NTP server fails.

The time server daemon synchronizes fabric time by sending updates of the Principal or Primary FCS local switch time periodically to every switch in the fabric. The time server daemon runs on all switches in the fabric, but only the Principal switch (when an FCS policy is not enabled) or the Primary FCS switch (when an FCS policy is enabled) connect to the NTP server (if specified) and broadcasts time service updates.

All switches in the fabric maintain the current clock server IP address in nonvolatile memory. By default, this value is **LOCL**, that is, the local clock of the Principal or the Primary FCS switch is the default clock server. Changes to the clock server IP addresses on the Principal or Primary FCS switch are propagated to all switches in the fabric.

Use this command without parameters to display the active NTP server and the configured NTP server list. Specify the **ipaddr** operands to set the clock server IP addresses and enable fabric-wide clock synchronization with the specified clock server. A value of **LOCL** may be specified as operand to set the clock server back to default.

Each **ipaddr** specified should be the IP address of an NTP server and should be accessible from the switch. The NTP server must support a full NTP client. When a clock server IP address other than **LOCL** is specified but is not used by the fabric, a warning is displayed and logged. When a clock server IP address other than **LOCL** is specified, the **date** command is restricted to display only. Refer to the **date** command help for details.

Notes When an FCS policy is enabled, this command can be run on all switches to view the clock server value. However, you can only modify the clock server value from the Primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

ipaddr

Specify the IP address of the NTP server or **LOCL** to use the local clock of the Principal or Primary FCS switch as the clock server. If more than one IP address is specified, separate the IP addresses by semicolons and enclose the list in double quotation marks. This operand is optional; if omitted, the current NTP server configuration is displayed. The default NTP server is **LOCL**.

Examples To display the default clock server:

```
switch:admin> tsclockserver
LOCL
```

tsClockServer

To set the NTP server to a specified IP address:

```
switch:admin> tsclockserver 123.123.123.123  
Updating Clock Server configuration...done.
```

```
switch:admin> tsclockserver  
123.123.123.123
```

To configure multiple NTP servers:

```
switch:admin> tsclockserver "12.134.125.24; 12.234.87.01"  
Updating Clock Server configuration...done.
```

SEE ALSO [date](#)

tsTimeZone

Displays or sets the system time zone.

Synopsis

tsTimeZone --interactive

tsTimeZone *timezonename*

tsTimeZone --old *hourOffset[,minuteOffset]*

Description

Use this command to display or set the system time zone.

All switches maintain the current time zone setup in nonvolatile memory. Changing the time zone on a switch updates the local time zone setup and is reflected in local time calculations.

All switches are by default in the 0,0 time zone:, which is, GMT. If all switches in a fabric are located in the same time zone, you may leave the default time zone setup.

Time zone is used in computing local time for error reporting and logging. An incorrect time zone setup does not affect the switch operation in any way.

System services started during the switch boot reflect a time zone change only at the next reboot.

The time zone can be specified in the following two ways, by name or in an hours and minutes offset format:

- The offset format is specified with the **--old** option, followed by an hour offset value and optionally a minute offset value.
- The time zone name format uses the timezone database, which automatically adjusts for Daylight Saving Time.

By default, the switch is in offset mode (**--old**), with zero offsets, that is, time is displayed in GMT. Use **tsTimeZone *timezonename*** to change the offset format to the timezone name format.

When executed without parameters, this command displays the current time zone configuration in the format in which it was set.

- Negative hour offset values mean that the local time is behind GMT; for example, -8,0 is GMT-08:00.
- Positive hour offset values mean the that local time is ahead of GMT; for example, 3,0 is GMT+03:00.

When Virtual Fabrics are enabled, the hardware clock is updated by the default switch in the chassis, and the time zone set on any logical switch applies to all logical switches on the chassis. The **tsTimeZone** command requires chassis permissions.

Because there is only one clock on the chassis, for the time server to function correctly, ensure that all logical switches in the fabric have the same NTP Clock Server configured. This includes any Pre-v6.2.0 switches in the fabric.

Notes

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

The option **11** (none) in interactive command mode to specify the time zone using the POSIX TZ format is not supported.

Operands This command has the following operands:

--old

Specifies the time zone in the offset format.

hourOffset

Specifies the number of hours relative to GMT. This operand must be specified as an integer. Valid values are -12 to 12. This operand is required with the **--old** option.

minuteOffset

Specifies the number of minutes relative to hour offset. This operand must be specified as an integer and is valid only with the **--old** option. Valid values are -30, 0, or 30. This operand is optional; if not specified, the value defaults to 0.

timezonename

Specifies the name of a time zone from the time zone database. Use **tstimezone --interactive** for a listing of valid time zone name.

--interactive

Interactively sets the timezone in name format.

Examples To display the current time zone setup and then change it to GMT-3:30:

```
switch:admin> tstimezone
Time Zone Hour Offset: 0
Time Zone Minute Offset: 0

switch:admin> tstimezone -3, -30
Updating Time Zone configuration...done.
System Time Zone change will take effect at next reboot.

switch:admin> tstimezone
Time Zone Hour Offset: -3
Time Zone Minute Offset: -30
```

To set the current timezone to the zone to Pacific Time using the interactive command mode:

```
switch:admin> tsTimeZone --interactive
Please identify a location so that time zone rules can be set correctly.
Please select a continent or ocean.
1) Africa
2) Americas
3) Antarctica
4) Arctic Ocean
5) Asia
6) Atlantic Ocean
7) Australia
8) Europe
9) Indian Ocean
10) Pacific Ocean
11) none - I want to specify the time zone using the POSIX TZ format.
Enter number or control-D to quit ?2
Please select a country.
1) Anguilla
2) Antigua & Barbuda
3) Argentina
4) Aruba
5) Bahamas
28) Haiti
29) Honduras
30) Jamaica
31) Martinique
32) Mexico
```

- | | |
|----------------------------------|-----------------------------|
| 6) Barbados | 33) Montserrat |
| 7) Belize | 34) Nicaragua |
| 8) Bolivia | 35) Panama |
| 9) Bonaire Sint Eustatius & Saba | 36) Paraguay |
| 10) Brazil | 37) Peru |
| 11) Canada | 38) Puerto Rico |
| 12) Cayman Islands | 39) Sint Maarten |
| 13) Chile | 40) St Barthelemy |
| 14) Colombia | 41) St Kitts & Nevis |
| 15) Costa Rica | 42) St Lucia |
| 16) Cuba | 43) St Martin (French part) |
| 17) Curacao | 44) St Pierre & Miquelon |
| 18) Dominica | 45) St Vincent |
| 19) Dominican Republic | 46) Suriname |
| 20) Ecuador | 47) Trinidad & Tobago |
| 21) El Salvador | 48) Turks & Caicos Is |
| 22) French Guiana | 49) United States |
| 23) Greenland | 50) Uruguay |
| 24) Grenada | 51) Venezuela |
| 25) Guadeloupe | 52) Virgin Islands (UK) |
| 26) Guatemala | 53) Virgin Islands (US) |
| 27) Guyana | |

Enter number or control-D to quit ?49

Please select one of the following time zone regions.

- 1) Eastern Time
- 2) Eastern Time - Michigan - most locations
- 3) Eastern Time - Kentucky - Louisville area
- 4) Eastern Time - Kentucky - Wayne County
- 5) Eastern Time - Indiana - most locations
- 6) Eastern Time - Indiana - Daviess, Dubois, \
 Knox & Martin Counties
- 7) Eastern Time - Indiana - Pulaski County
- 8) Eastern Time - Indiana - Crawford County
- 9) Eastern Time - Indiana - Pike County
- 10) Eastern Time - Indiana - Switzerland County
- 11) Central Time
- 12) Central Time - Indiana - Perry County
- 13) Central Time - Indiana - Starke County
- 14) Central Time - Michigan - Dickinson, Gogebic, \
 Iron & Menominee Counties
- 15) Central Time - North Dakota - Oliver County
- 16) Central Time - North Dakota - Morton County \
 (except Mandan area)
- 17) Central Time - North Dakota - Mercer County
- 18) Mountain Time
- 19) Mountain Time - south Idaho & east Oregon
- 20) Mountain Time - Navajo
- 21) Mountain Standard Time - Arizona
- 22) Pacific Time
- 23) Alaska Time
- 24) Alaska Time - Alaska panhandle
- 25) Alaska Time - southeast Alaska panhandle
- 26) Alaska Time - Alaska panhandle neck
- 27) Alaska Time - west Alaska
- 28) Aleutian Islands
- 29) Metlakatla Time - Annette Island
- 30) Hawaii

Enter number or control-D to quit ?18

The following information has been given:

tsTimeZone

```
United States
Mountain Time
Therefore TZ='America/Denver' will be used.
Local time is now: Tue Nov 6 02:43:16 MST 2012.
Universal Time is now: Tue Nov 6 09:43:16 UTC 2012.
Is the above information OK?
1) Yes
2) No
Enter number or control-D to quit ?1
System Time Zone change will take effect at next reboot
```

To revert back to the offset format and verify the configuration:

```
switch admin> ts timezone --old 2

switch admin> ts timezone
Time Zone Hour Offset: 2
Time Zone Minute Offset: 0
```

SEE ALSO [date](#)

turboRamTest

Performs a turbo SRAM test of ASIC chips.

Synopsis `turboramtest`
 `[--slot slot]`
 `[-passcnt count]`

Description Use this command to verify the chip SRAM located in the ASIC using the turbo-RAM BIST circuitry. The BIST controller is able to perform the SRAM write and read operation at a much faster rate than the PCI operation.

Notes Do not abort this test prematurely, using **CTRL-C** or **q** to quit. Doing so may cause the test to report unexpected errors. Errors may vary depending on the hardware platform.

You cannot interrupt the test by pressing the return key (<cr>).

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This diagnostic cannot be run on an operational switch. You must disable the switch using the **chassisDisable** command before you can run this test. After the test completes, re-enable the switch using the **chassisEnable** command.

Operands This command has the following optional operands:

--slot *slot* Specifies the slot number on which the diagnostic operates. All eligible blade ports in the specified slot are tested. This operand is optional. The default value is 0 and operates on fixed port count products.

-passcnt *count* Specifies the number of times to perform this test. This operand is optional. The default value is 1.

Examples To run the SRAM test with two passes:

```
switch:admin> turboramtest -passcnt 2
Platform bladeType 154

Running turboramtest .....

WARNING:
This test should NOT be aborted in the middle. If aborted, current
blade or fixed configuration switch may become unusable.
```

turboRamTest

Reset the blade or the switch to recover.

```
Running Condor3 turboramtest on slot 0 Condor 0  
Condor3 turboramtest on slot 0 condor 0 PASSED
```

PASSED.

SEE ALSO [bladeDisable](#), [bladeEnable](#), [chassisDisable](#), [chassisEnable](#)

upTime

Displays length of time the system has been operational.

Synopsis **uptime**

Description This command displays the current time, how long the system has been running, how many users are currently logged on, and the system load averages for the past 1, 5, and 15 minutes.

If the uptime is less than 60 seconds, the time is displayed in seconds. For times greater than or equal to 60 seconds, the time is displayed in minutes. The output format adjusts accordingly.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display the length of time the system has been operational:

```
switch:admin> uptime
12:03am up 4:56, 3 users, load average: 1.17, 1.08, 1.08
```

SEE ALSO [date](#), [fastBoot](#), [reboot](#)

uRouteShow

Displays unicast routing information.

Synopsis `urouteshow [slot][port] [domain]`

Description Use this command to display the unicast routing information for a port, as it is known by the FSPF path selection and routing task. The routing information describes how a frame that is received from a port on the local switch is to be routed to reach a destination switch.

The following information displays:

Local Domain ID

Domain number of local switch.

In Port

Port from which a frame is received. Except for the cases in which you perform a port swap or enable extended-edge PID (PID Format 2) on a switch, the value is equal to the port index field displayed by the **switchShow** command. Refer to *Fabric OS Administrator's Guide* for more information regarding the extended edge PID format.

Domain

Destination domain of incoming frame.

Out Port

Port to which the incoming frame is to be forwarded. Except for the cases in which you perform a port swap or enable extended edge PID (PID Format 2) on a switch, the value is equal to the port index field displayed by the **switchShow** command. For port swap operations, the value is equal to the switch port field displayed by the **portSwapShow** command. Refer to *Fabric OS Administrator's Guide* for more information regarding the extended-edge PID format.

Metric

Cost of reaching the destination domain.

Hops

Maximum number of hops required to reach the destination domain. If the number of hops are different for some multiple equal cost paths (to reach the same domain), then it displays the maximum number.

Flags

Indicates rout type as either dynamic (D) or static (S). A dynamic route is discovered automatically by the FSPF path selection protocol. A static route is assigned using the command **uRouteConfig**.

Next (Dom, Port)

Domain and port number of the next hop. These are the domain number and the port number of the switch to which Out Port is connected.

The information provided by this command should match what is provided by **portRouteShow** and **topologyShow**.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

slot

For bladed systems only, specify the slot number of the input port whose routes are displayed, followed by a slash (/).

port

Specify the number of the input port whose routes are to be displayed, relative to its slot for bladed systems. Use **switchShow** to display a list of valid ports. This operand is optional; if omitted, the command displays routing information for all input ports in the switch.

domain

Specify a remote domain in the fabric for which routing information is to be displayed. This operand is optional; if omitted, the routing information for all domains in the fabric is displayed.

Examples To display the routing information of all the active ports:

```
switch:admin> urouteshow
Local Domain ID: 3
In Port Domain Out Port Metric Hops Flags Next (Dom, Port)
-----
  0         1      11    1000   1   D           1,0

Type <CR> to continue, Q<CR> to stop:
11         2       0    1500   2   D           4,0

Type <CR> to continue, Q<CR> to stop:
         4       0     500   1   D           4,0
16         1      27    1000   1   D           1,1

Type <CR> to continue, Q<CR> to stop:
27         2      16    1500   2   D           4,16
```

To display the routing information of port 11 to domain 4 only:

```
switch:admin> urouteshow 1/11, 4
Local Domain ID: 3
In Port Domain Out Port Metric Hops Flags Next (Dom, Port)
-----
  11         4      16     500   1   D           4,16
```

SEE ALSO [portRouteShow](#), [topologyShow](#)

usbStorage

Manages data files on an attached USB storage device.

Synopsis

```
usbstorage [-e | --enable]
usbstorage [-d | --disable]
usbstorage [-l | --list]
usbstorage [-r | --remove area target]
usbstorage [-h | --help]
```

Description Use this command to control a USB device attached to the Active CP. When the USB device is enabled, other applications, such as **supportSave**, **firmwareDownload**, or **configDownload/cfgUpload** can conveniently store and retrieve data from the attached storage device. Refer to the help pages for these commands for specific information on how the USB device is accessed by each application.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

This command is available only on the Active CP.

On the Brocade 7840, the USB may continue blinking even after the USB device is disabled.

Operands This command has the following operands:

-e | --enable Enables the USB device. The USB device must be enabled before the list and remove functions are available.

-d | --disable Disables an enabled USB device. This command prevents access to the device until it is enabled again.

-r | --remove *area target* Removes a target in a specified application storage area. Valid areas are: *firmware*, *support*, *config*, or *firmwarekey*. A specified *area* must be followed by a *target*. Any existing file in a given application directory can be specified as a target. This command removes all data associated with the specified target.

-l | --list Lists the content of the USB device up to two levels down from the root directory.

-h | --help Displays the command usage.

Examples To enable an attached USB device:

```
switch:admin> usbstorage -e
USB storage enabled
```

To list the contents of the attached USB device:

```
switch:admin> usbstorage -l

firmwarekey\          0B      2010 Aug 15 15:13
support\              106MB   2010 Aug 24 05:36
  support1034\        105MB   2010 Aug 23 06:11
config\               0B      2010 Aug 15 15:13
firmware\             380MB   2010 Aug 15 15:13
  FW_v6.4.0\         380MB   2010 Aug 15 15:13
  Available space on usbstorage 74%
```

To remove a firmware target from the firmware application area:

```
switch:admin> usbstorage -r firmware FW_v6.0.0
```

To disable an attached USB device:

```
switch:admin> usbstorage -d
USB storage disable
```

SEE ALSO [supportSave](#), [firmwareDownload](#), [configUpload](#), [configDownload](#)

userConfig

Manages user accounts.

Synopsis

```

userconfig

userconfig --show [username | -a | -r rolename]

userconfig --showlf -l LF_ID | -c

userconfig --add username -r role -l LF_ID_list [-h LF_ID
  [-c chassis_role] [-d description] [-p password]
  [-at | -access-time start_time-end_time] [-x]

userconfig --change username [-r role] [-h LF_ID]
  [-l LF_ID_list] [-c chassis_role] [-d description] [-x]
  [-e yes | no] [-u]

userconfig --addlf username [-h LF_ID]
  [-r role -l LF_ID_list] [-c chassis_role]

userconfig --deletelf username [-h LF_ID][-l LF_ID_list] [-c]

userconfig --delete username

userconfig --help

```

Description

Use this command to manage user accounts on a switch. In a Virtual Fabric-enabled environment, you can configure the account's username, its role, and the logical fabrics that the account may access. An account can have different roles for different Logical Fabrics. An account can access multiple Logical Fabrics, but only one Logical Fabric at a time.

When executed without operands, this command displays the usage. The logical fabric command options are displayed only if Logical Fabrics are enabled on the switch.

You can execute the **userConfig** command on any switch. When the command completes, account information is saved persistently. On platforms supporting multiple switch domains, account information is saved only to the switch domain, in which the command was executed.

Use the **distribute** command to distribute the account database manually to other switches in the fabric. Target switches must be configured to accept the database. Accounts that are not consistent with the distributed database are overwritten. Account recovery from backup or access to backup data is not supported..

This command supports all user-defined roles in addition to the default roles provided with Fabric OS. To display the user-defined roles with their associated access permissions, use the **roleConfig --show -all** command. To display all roles, including the default roles, use the **roleConfig --show -all -default** command.

In a Logical Fabric environment, you can additionally define access to chassis-level commands. An account can have one role in the Logical Fabric, and another role regarding chassis commands.

An asterisk (*) next to the account name in the **userConfig --show** and **userConfig -a** output indicates that the password for that user account is still at the default value.

Notes The **userConfig** command operates on the switch-local user database only, regardless of whether the switch is configured to use RADIUS or LDAP authentication.

The account database supports a maximum of 256 customer created accounts.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

--show

Displays user account information. Only accounts with access permissions compatible with the SecurityAdmin and Admin roles can show information about accounts other than the current login account. The following operands are optional:

username

Specifies the account login name. When no operand is specified, the command displays the current account information.

-a

Displays information about all accounts.

-r rolename

Displays information about all accounts with the specified role.

--showlf

Displays Logical Fabric permissions in an Logical Fabrics-enabled environment. Only users with access permissions compatible with the SecurityAdmin or Admin role may execute this command. An operand is required with this command. The following operands are mutually exclusive:

-l *LF_ID_list*

For each Logical Fabric in the *LF_ID_list*, this option displays a list of users that include that Logical Fabric in their Logical Fabrics permissions. Specify a range (1-5), or a list of Logical Fabric IDs separated by a comma (1,2,3), or a combination of both (1-5,7). Only users with access permissions compatible with the SecurityAdmin or Admin role may execute this command.

-c

Displays a list of users who have permission to execute chassis commands.

--add | --change

Creates a new user account or modifies an existing user account. The following restrictions apply when you create or modify a user account:

- You cannot change the role, the Logical Fabric permissions, the home Logical Fabric of any default account.
- You cannot change the role, the Logical Fabric permissions, or the description of accounts at the same or a higher authorization level.
- You cannot change the role, the Logical Fabric permissions, or the home Logical Fabric of your own account.
- No account can disable itself.
- Logical Fabric permissions must be a subset of the respective Logical Fabric permissions of the account that creates or modifies a user account.
- In an Logical Fabric-enabled environment, you can change the role associated with existing Logical Fabrics but you cannot add new Logical Fabrics or delete any existing Logical Fabrics.
- You cannot use **change** if the default FID was modified after user creation. **--addlf** must be used to add newly created Logical Fabrics to user.
- The account name cannot be the same as an existing user account, a default role, a user-defined role, or a system role. System roles are used by internal switch processes and include the following: bin, daemon, sys, adm, tty, disk, lp, mem, kmem, wheel, mail, news, uucp, man, dip, ftp, nobody, users, floppy, console, utmp, and slocate. If the specified username already exists, this command fails with an appropriate message. Choose a different username and reissue the command.

The following operands are supported with the **--add** and **--change** option

username

Specifies the login name of the account to be created or modified. Enter a valid login name to modify an existing account. For new accounts, the name must be unique and must begin with an alphabetic character. User names are case-sensitive and can contain up to 32 alphanumeric characters, including periods (.) and underscore (_) characters. For backward compatibility with releases prior to Fabric OS v7.3.0, the username with 40 characters is supported through configDownload and password distribution from switches.

-r role

In an Logical Fabric-enabled environment, this operand specifies the account's role for all Logical Fabrics provided with the Logical Fabric list. When you create a user account in an Logical Fabric-enabled environment, you can specify only one role for the user. This role is associated with each of the Logical Fabric IDs in the specified LF_ID_list. Once the account is created, you can use the **--addlf** option to create another list of Logical Fabric IDs with its own set of associated roles.

You can assign any role to the account, user-defined or default. Use the **roleconfig --show** command for a listing of valid roles. This operand is required with the **--add** option; it is optional with the **--change** option.

-h LF_ID

Specifies the home Logical Fabric depending on the environment. This operand is optional. If no Logical Fabric is specified with the **--add** option, the system assigns the lowest numbered Logical Fabric the user is authorized to access.

- l *LF_ID_list***
Specifies the Virtual Fabrics the user is authorized to access. The Logical Fabrics in *LF_ID_list* and the existing Logical Fabric permissions for username must be a subset of the Logical Fabric permissions of the account that executes this command. This operand is required with the **--add** option. It is optional with the **--change** option.
- c *chassis_role***
Specifies the account's access permissions at the chassis level. The chassis role allows the user to execute chassis-related commands in an Logical Fabric-enabled environment. To assign the chassis role to an account, the executing account must have chassis-level permissions. Valid chassis roles include the default roles and any of the user-defined roles. Use the **roleconfig --show** command for a listing of valid roles.
- d *description***
Provides a description for the new account. This operand is optional. The maximum length is 32 printable ASCII characters. Some characters that are interpreted by the shell (" , ; ! etc.) require a preceding escape character (\). To include spaces, place the description in double quotation marks. Colons are not permitted.
- at | -access-time *start_time-end_time***
Specifies the time range the users can access the switch through Telnet, SSH, console or Web. Outside the configured time, access will be denied and the existing sessions for the user will be terminated on the configured *end_time*. The *start_time* and *end_time* must be specified in the *hh:mm* 24-hour clock format.

If *end_time* is less than the *start_time*, then it is considered as next day end time. For example, the command **-access-time 23:20-07:20** represents time interval in between two days.
- x**
Optionally specifies an expired password that must be changed the first time the user logs into a new or modified account. This command also prompts for the existing password.

The following optional operand is available only with the **--add** option:
- p *password***
Specifies a password for the account. This operand is optional; if omitted, the command prompts for an initial password for the account. The password must satisfy all currently enforced password rules. By default the password is created with the configured expiration period.

The following optional operands are available only with the **--change** option:
- e yes | no**
Enables or disables an account. Specify "yes" to enable or "no" to disable an account. Once an account is disabled, the CLI sessions associated with the account are terminated.
- u**
Unlocks the specified user account. User accounts can get locked after several attempts to log in with an invalid password. Refer to the **passwdCfg** help page for more information.

--addlf | --deletelf

Adds Logical Fabric or chassis permissions to a user account or deletes Logical Fabric or chassis permission from a user account. The following operands are supported:

username

Specifies the account login name.

-h *LF_ID*

Specifies the account's home Logical Fabric. This operand is optional.

- If a home Logical Fabric is specified with the **--addlf** option, the home Logical Fabric must be one of the Logical Fabrics in *LF_ID_list*. If a home Logical Fabric is not specified and the account did not previously have a home Logical Fabric, the home Logical Fabric is set to the lowest numbered Logical Fabric in the user's Logical Fabric permissions.
- If a home Logical Fabric is specified with the **--deletelf** option, the home Logical Fabric must be one of the Logical Fabrics in the Logical Fabric permissions remaining after the Logical Fabrics specified in *LF_ID_list* have been removed. If a home Logical Fabric is not specified, the current home Logical Fabric remains unchanged, if it is still in the user's Logical Fabric permissions. If a home Logical Fabric is not specified and the current home Logical Fabric is deleted, the new home Logical Fabric is set to the lowest numbered Logical Fabric in the user's Logical Fabric permissions.

The account's existing Logical Fabric permission and the *LF_ID_list* must be a subset of the Logical Fabric permissions of the account executing this command.

-l *LF_ID_list*

Specifies the logical fabric IDs to be added or deleted. Specify a range (1-5) or a list of Logical Fabric IDs separated by comma (1,2,3), or a combination of both (1-5,7).

-r *role*

Specifies the role associated with the Logical Fabric ID list given in this command. This operand is required when you specify an *LF_ID_list* operand.

-c [*chassis_role*]

Specifies the account's access permissions regarding chassis-level commands. To remove an account's chassis permissions, specify **-c** only. To add chassis permissions, specify a chassis role with the **--c** option.

--delete username

Deletes the specified account from the switch. This command prompts for confirmation. Once an account is deleted, the CLI sessions associated with the account are terminated.

The following restrictions apply when you delete an account:

- You cannot delete a default account.
- You cannot delete your own account.

LF_ID list and associated Logical Fabric permissions for *username* must be a subset of the Logical Fabric permissions of the account that executes the **userConfig --delete** command.

--help

Displays the command usage. In a Logical Fabric-enabled environment, options specific to Logical Fabrics are displayed.

Examples To create a new account named "test" with admin role and admin chassis permissions in the Logical Fabric member list 1-10:

```
switch:admin> userconfig --add test -l 1-10 \
-r admin -c admin
Setting initial password for test
Enter new password:
Re-type new password:
Account test has been successfully added.
```

To display current account information:

```
switch:admin> userconfig --show test
Account name: test
Role: admin
Description:
Enabled: Yes
Password Last Change Date: Sat Jun 14 2008
Password Expiration Date: Not Applicable
Locked: No
RoleLFMaps: admin: 1-10 chassis
Chassis Role: admin
Home Context: 1
```

To grant user access permissions to the test account for the Virtual Fabrics 11-15:

```
switch:admin> userconfig --addlf test -r user -l 11-15
New LFs/Chassis role for account test has been \
successfully added.
```

To change the test account's access permissions for the Logical I Fabrics 5 and 6 to ZoneAdmin and the chassis permission to user:

```
switch:admin> userconfig --change test -r zoneadmin \
-l 1-5 -c user -h 4
Broadcast message from root (ttyS0) Sat Jun 14 01:05:28 2008...
Security Policy, Password or Account Attribute Change: \
test will be logged out
```

To display the test account information:

```
switch:admin> userconfig --show test
Account name: test
Role: zoneadmin
Description:
Enabled: Yes
Password Last Change Date: Sat Jun 14 2008
Password Expiration Date: Not Applicable
Locked: No
RoleLFMaps: zoneadmin: 1-5 admin: 6-10 user: 11-15 chassis
Chassis Role: user
Home Context: 4
```

userConfig

To remove chassis permissions from the test account for the Logical Fabrics 1-3.

```
switch:admin> userconfig --deletelf test -l 1-3 -c
Broadcast message from root (ttyS0) Sat \
Jun 14 01:10:02 2008...

Security Policy, Password or Account Attribute Change: \
test will be logged out
LFs/chassis role for account test has been successfully\
deleted.
```

To display information for all accounts with admin privileges:

```
switch:admin> userconfig --show -r admin

Account name: admin
Description: Administrator
Enabled: Yes
Password Last Change Date: Unknown
Password Expiration Date: Not Applicable
Locked: No
Home LF Role: admin
Role-LF List: admin: 1-128
Chassis Role: admin
Home LF: 128

Account name: test1s1
Description:
Enabled: Yes
Password Last Change Date: Sun Oct 5 2025
Password Expiration Date: Not Applicable
Locked: No
Home LF Role: admin
Role-LF List: admin: 1
No chassis permission
Home LF: 1
```

SEE ALSO [roleConfig](#)

version

Displays firmware version information.

Synopsis `version`

Description Use this command to display firmware version information and build dates.

The command output includes the following:

Kernel The version of switch kernel operating system.

Fabric OS The version of switch Fabric OS.

Made on The build date of firmware running in switch.

Flash The build date of firmware stored in flash prompts.

BootProm The version of the firmware stored in the boot PROM

Usually the Made on and Flash dates are the same, because the switch starts running flash firmware at power-on. However, in the time period between **firmwareDownload** and the next **reboot**, the dates can differ.

Operands None

Examples To display the firmware version information in a switch:

```
switch:admin> version
Kernel:      2.6.14.2
Fabric OS:   v7.2.0
Made on:     Wed Apr 24 19:52:07 2013
Flash:       Fri Apr 26 06:58:33 2013
BootProm:    1.0.15
```

SEE ALSO [firmwareDownload](#), [reboot](#)

wwn

Displays the world wide name (WWN) and factory serial number of the switch or chassis.

Synopsis `wwn [-sn]`

Description Use this command to display the WWN associated with a switch or chassis and to display the factory serial number. The WWN is a 64-bit number that has eight colon-separated fields each consisting of one or two hexadecimal digits between 0 and ff. The WWN is a factory-set parameter that cannot be changed by the end user. The WWN is used as the license ID in many cases, but the only official string to be used for requesting licenses is the **licenseidShow** output. Alternately, use **switchShow** to display the WWN.

In addition to the WWN, all switches have a unique 24-bit Fibre Channel address that is used for communicating with the switch. Use **fabricShow** to display the FC address in addition to the WWN.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operands:

-sn

On enterprise-class platforms, this operand displays the chassis factory serial number following the WWN. On standalone platforms, it displays the factory serial number. This operand is optional; if omitted, this command displays only the WWN for the switch or chassis.

Examples To display the WWN on a Brocade DCX-4S:

```
switch:admin> wwn
10:00:00:05:1e:7a:7a:00
```

To display the WWN and chassis factory serial number:

```
switch:admin> wwn -sn

WWN: 10:00:00:05:1e:7a:7a:00
SN:  ANP0645D05B

switch:admin> chassisshow | grep ANP0645D05B
Chassis Factory Serial Num: ANP0645D05B
switch:admin>
```

To display the license ID:

```
switch:admin> licenseidshow
10:00:00:05:1e:7a:7a:00
```

To display the WWN and the Fibre Channel address:

```
switch:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr FC IP Addr Name
-----
1:fffc01 10:00:00:05:1e:7a:7a:00 10.32.39.25 0.0.0.0 "sw03"
2:fffc02 10:00:00:05:1e:b3:00:9e 10.32.39.34 0.0.0.0 "ras39"
3:fffc03 10:00:00:05:1e:93:c4:00 10.32.39.20 0.0.0.0 "sw5"
4:fffc04 10:00:00:05:1e:55:5c:69 10.32.39.59 0.0.0.0 >"sp39"
```

To display the WWN on a Brocade 5100:

```
switch:admin> wwn
10:00:00:05:1e:7a:7a:00
```

To display the WWN and factory serial number:

```
switch:admin> wwn -sn

WWN: 10:00:00:05:1e:82:3c:2a
SN:  ALM0602E003

switch:admin> chassisshow | grep ALM0602E003
Factory Serial Num:      ALM0602E003
```

SEE ALSO [chassisShow](#), [fabricShow](#), [licenseIdShow](#), [switchShow](#)

wwnAddress

Binds an FC Port ID to a device WWN.

Synopsis `wwnaddress --bind [WWN] [PID]`

`wwnaddress --unbind [WWN]`

`wwnaddress --show`

`wwnaddress --findPID [WWN]`

`wwnaddress --help`

Description

Use this command to manage address assignments for a given device world wide name. The allocation of a PID to a specified device WWN supports the persistence of the PID based on the WWN of the device to which the PID is bound. If the PID is not bound to a device WWN, the device can get the same or a different PID irrespective of which port it logs in to a given switch partition.

This command fails under any of the following conditions:

- The PID is currently bound to another port through port address binding. Use `portaddress --unbind` to free up the PID.
- The WWN is already bound with a different PID, or the PID is bound to another WWN. Use `wwnaddress --unbind` to remove the PID-WWN binding.
- There is no space left in the WWN-PID table for an additional entry. Use `wwnaddress --unbind` to free up space in the table.

The command provides a `--show` option that displays the currently bound address for all WWNs. Alternately, you can use the `--findPID` option to display the PID currently bound to a specified WWN.

Notes

Dynamic Area Mode and WWN-Based persistent PID must be enabled on the switch before you can assign an address with this command. Refer to **configure** for more information.

This command is supported only on the Brocade DCX and on the DCX-4S.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands

This command has the following operands:

`--bind`

Assigns the lower two bytes of the Fibre Channel address to the specified WWN.

`WWN`

Specifies the WWN for the device to which the PID should be assigned.

`PID`

Specifies the PID (the lower 16 bits of the address excluding the domain part) in hexadecimal format to be bound to the device WWN. Note that only the upper 10 bits of the PID can be used for a unique route. Therefore, not all addresses in the 16-bit range are available.

| | |
|------------------|--|
| --unbind | Removes the WWN-PID binding corresponding to the specified device WWN. |
| WWN | Specifies the WWN for the device from which to remove the PID binding. |
| --show | Displays all WWN-PID entries currently present in the partition. |
| --findPID | Displays the PID currently bound to the specified device WWN. |
| WWN | Specifies the device WWN. |
| --help | Displays the command usage. |

Examples To bind a 16-bit address to a device WWN:

```
switch:admin> wwnaddress --bind \
10:00:00:06:2b:0f:76:5f 0x9000
```

To unbind a given address from a WWN:

```
switch:admin> wwnaddress --unbind \
10:00:00:06:2b:0f:76:5f
```

To display all WWN address bindings on the current partition:

```
switch:admin> wwnaddress --show
#   WWN                               Area Age Flag
=====
1)  10:00:00:06:2b:0f:71:0c             0x405  53  0x12
2)  10:00:00:05:1e:5e:2c:11             0x9900 101 0x21
3)  10:00:00:06:2b:0f:71:0d             0x503  37  0x12
4)  10:00:00:06:2b:0f:71:0e             0x304  43  0x12
5)  10:00:00:06:2b:0f:71:0f             0x303  38  0x12
6)  10:00:00:06:2b:0f:70:14             0x401  29  0x12
7)  10:00:00:06:2b:0f:70:15             0x505  46  0x12
8)  10:00:00:06:2b:0f:70:16             0x402  33  0x12
9)  10:00:00:06:2b:0f:70:17             0x406  47  0x12
10) 10:00:00:06:2b:0f:72:20              0x403  30  0x12
11) 10:00:00:06:2b:0f:72:21              0x501  36  0x12
12) 10:00:00:06:2b:0f:72:23              0x502  34  0x12
13) 10:00:00:06:2b:0f:6e:30              0x301  35  0x12
14) 10:00:00:06:2b:0f:6e:31              0x102  42  0x12
15) 10:00:00:06:2b:0f:6e:32              0x302  39  0x12
16) 10:00:00:06:2b:0f:6e:33              0x504  45  0x12
17) 10:00:00:06:2b:0f:76:5e              0x404 101 0x12
18) 10:00:00:06:2b:0f:76:5f              0x101  41  0x12
19) 20:20:00:05:1e:0b:61:cc              0x400  28  0x22
20) 20:21:00:05:1e:0b:61:cc              0x500  31  0x22
21) 20:22:00:05:1e:0b:61:cc              0x300  32  0x22
22) 20:23:00:05:1e:0b:61:cc              0x100  40  0x22
23) 10:00:00:06:2b:0f:6d:ee              0x305  50  0x12
24) 10:00:00:06:2b:0f:6d:ef              0x103  49  0x12
```

wwnAddress

To display the WWN address binding for device 20:22:00:05:1e:0b:61:cc:

```
switch:admin> wwnaddress --findPid 20:22:00:05:1e:0b:61:cc
WWN                               PID
=====
20:22:00:05:1e:0b:61:cc  0x300
```

SEE ALSO [portAddress](#)

wwnRecover

Utility to recover mismatch of data in WWN card. Errors may be due to corruption, defect, or other causes.

Synopsis `wwnrecover`

Description Use this command to recover one or both WWN cards from errors detected by the audit routine. The audit routine is triggered upon first boot up of the system, periodically every 24 hours thereafter, and upon detection of insertion of either of the WWN card. The audit routine performs consistency checks on each of the following sections of data maintained in the WWN cards for use by the corresponding system component:

- License ID data sections within the same WWN card
- IP address data sections across the two WWN cards
- DCE data sections across the two WWN cards
- License ID data sections across the two WWN cards
- Chassis configuration data sections across the two WWN cards
- Chassis serial and part number sections across the two WWN cards

Error resulting from a mismatch in data may be encountered on one or more sections. The recovery mechanism may vary depending on the nature of the error encountered and is also based on the data section. The **wwnrecover** operation presents a summary of all the errors detected and prompts you through the recovery process interactively for all data sections. For each data section, you will be prompted with an option to recover WWN1 card data from WWN2 card data or vice-versa or to exit without recovering the corresponding data section. For example, if WWN2 is replaced by a minimally programmed factory replacement card, WWN recovery recommends to recover most of the data sections on WWN2 from WWN1. However, in a rare case where a bit corruption occurs in one data section on one card and another bit corruption in a different data section on the other card, the recovery for each of the affected data sections will be different. WWN recovery is not possible if the license ID differs on both WWN cards.

In the event of any of the following issues with WWN card operations, the system logs the corresponding RASlog message:

- Mismatch between the two WWN cards detected by the audit routine (EM-1220 message).
- Insertion of either of the two WWN cards detected by the system (EM-1221 message).
- Access failure by the audit routine during read/write to either of the WWN cards (EM-1222 message).

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To perform WWN recovery (the following example shows that no WWN errors were detected):

```
switch:admin> wwnrecover
No WWN Error is detected. Exiting.
```

To perform WWN recovery (the following example shows the WWN errors and the recovery process):

```
switch:admin> wwnrecover
WWN Discrepancies (Error:64) detected.

Please attempt recovery of these errors (where possible)
by navigating to the required recovery selection
from the following WWN Recovery Options Menu.

Please note that in the event of a mismatch in the LicenseIds,
between WWN1 and WWN2 cards, wwn recovery may not be possible.
In this case a factory programmed card with a matching LicenseId
must be used, as a replacement card, to complete the wwn recovery.

WWN 1 and WWN 2 Non-Critical Seeprom data is mismatched.

      WWN Non-Critical Seeprom Problem Details

WWN Seeprom Chassis Serial Number Mismatch.
WWN 1 Serial Number: BADSerial
WWN 2 Serial Number: GOODSerial

      WWN Recovery Options

0. Exit

3. Recover WWN 2 from WWN 1

4. Recover WWN 1 from WWN 2

Enter Selection > 4
You have opted to recover WWN1 from WWN 2.
Are you sure? Please Confirm: (yes, y, no, n): [no] y
Recovering Serial Number...
Serial Number Recovery completed.

WWN Seeprom Chassis Part Number Mismatch.
WWN 1 part Number: GoodPARTNO
WWN 2 Part Number: BadPARTNo

      WWN Recovery Options

0. Exit

3. Recover WWN 2 from WWN 1

4. Recover WWN 1 from WWN 2

Enter Selection > 3
You have opted to recover WWN2 from WWN1.
Are you sure? Please Confirm: (yes, y, no, n): [no] y
Recovering Part Number...
Part Number Recovery completed.
```

SEE ALSO [licenseIdShow, wwn](#)

zone

Performs specific zone operations, manages Traffic Isolation (TI) Zones, and Frame Redirect (RD) Zones.

Synopsis

```

zone --help

zone --copy source_zone_object
      [dest_zone_object][-f]

zone --expunge "zone_object"

zone --validate [[-f]] [-m mode] ["zone_object"]
               [[-i] [pattern]]

zone --create -t objecttype [-o optlist] name -p portlist

zone --add [-o optlist] name -p portlist

zone --remove name -p portlist

zone --delete name

zone --activate name

zone --deactivate name

zone --show [-ic] [name][ -ascending]

zone --showTlerrors

zone --rdcreate [host_wwn] [target_wwn]
               [vi_wwn] [vt_wwn] [policy] [FCR | noFCR]

zone --rddelete name

zone --showTltrunkerrors

```

Description

The **zone** command supports three types of operations: specific zone operations, management of Traffic Isolation Zones, and management of Frame Redirect (RD) Zones.

1. Manage Zoning Operations

Use the **--copy**, **--expunge**, and **--validate** options to perform specific zoning operations. These commands follow a batched-transaction model.

2. Manage Traffic Isolation (TI) Zones

Use the **--create**, **--add**, **--remove**, **--delete**, **--activate**, **--deactivate**, and **--show** options to manage Traffic Isolation Zones.

TI zones control the flow of interswitch traffic by provisioning certain E_Ports to carry only traffic flowing from a specific set of source ports. The provision (a.k.a TI Zone) is part of the defined zone configuration and does not appear in the effective zone configuration. A Maximum of 255 TI Zones can be created in one fabric.

TI zones over FCR provide the ability to lock down a request and corresponding response to the same VE_Port tunnel for a given pair of devices in two separate fabrics. TI over FCR has two working parts:

- TI within edge fabric routes traffic between a real device and a Proxy device to a specified EX_Port.
- TI within backbone fabric locks down the route within the backbone fabric based on EX_Ports and devices involved.

Use the **--showTlerrors** option to generate a report of potential routing problems in the local Domain. If the command detects errors, it outputs the ID of the current domain, and for each record, it displays the following information:

Error Type

Error or Warning.

Affected Remote Domain

Domain where the traffic drop is likely to be happening.

Affected Local Port

Device port shared by the TI zones.

Affected TI Zones

Lists the names of the TI zones implicated in the problem.

Affected Remote Ports

Remote ports that are affected by the problem in the TI Zones.

Using TI zones in logical fabrics has several restrictions. TI zones created in logical fabrics cannot include logical interswitch links (LISLs); only physical port numbers can be included. TI zones cannot use the failover disabled option in Logical fabrics that have LISLs, and TI zones cannot use the failover disabled option in Base Fabrics. For more information, refer to the *Fabric OS Administrator's Guide*.

3. Manage Frame Redirect (RD) Zones

Use the **--rdcreate** and **--rddelete** options to manage Frame Redirect Zones. RD zones allow frames to be redirected to devices that can perform additional processing on these frames (for example, encryption). The feature uses a combination of RD zones and Name Server changes to map real device world wide names (WWNs) to virtual port IDs (PIDs.) This allows redirecting a flow between a host and target to a device that can perform its functionality without reconfiguring the host and target.

The RD Zone is part of the defined zone configuration and does not appear in the effective zone configuration. Use **cfgSave** to save the RD zone changes to the defined configuration. Use **cfgShow** to display the RD zones.

When you create the first RD zone, two additional zone objects are created automatically: A base zone "red_____base" and a zone configuration "r_e_d_i_r_c__fg". These additional zone objects are required by the implementation. These zone objects must remain on the switch as long as other redirect zones are defined. Do not remove these objects, unless redirection is no longer required and no other redirect zones exist.

Use **zone --rddelete** to remove the base RD zone, "red_____base", if it is the last remaining RD zone in the RD zone configuration, "r_e_d_i_r_c__fg." When the base zone is removed, the RD zone configuration is removed as well.

Notes Device ports are allowed to be part of several TI zones to support enhanced TI zone deployment in FICON environments.

You cannot swap E_Ports that are configured as part of a TI zone. The TI zone information is lost when you use **portSwap** to swap the E_Ports. To work around this issue, reconfigure your TI zones rather than swapping the E_Ports.

The current zone commands, **zoneCreate**, **aliCreate**, **cfgCreate**, etc., cannot be used to manage special zones, such as TI zones or RD zones.

Peer Zones are not allowed to be edited using this command.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command takes as an operand an action and its associated arguments. When executed without operands, the command displays the usage.

--help Displays the command usage.

1. Commands for performing specific zone operations

--copy Copies a specified zone object or all zone objects. The current transaction buffer is used for this operation. The following operands are optional:

source_zone_object
Specifies the zone object. The zone object can be a zone configuration, a zone alias, or several zones. If a source zone object is not specified, all zone configurations are copied over.

dest_zone_object
Identifies the destination zone object. If *dest_zone_object* is not specified, *source_zone_object* is copied over with the same name. If the destination zone object is not already present, one is created (with type as *source_zone_object*).

-f Overwrites existing zone object without confirmation.

--expunge Removes all references to the specified zone object and then deletes the zone object. The command displays the list of zone objects to be deleted and prompts for confirmation before deleting the zone objects. The removal of zone object references can trigger the removal of zones not specified in the command. For example, removing the last zone member from a zone deletes the zone, and may trigger a recursive deletion of other zones. The following operand is required:

"zone_object"
Specifies a zone object. A zone object can be a zone member, zone alias, or a zone. The zone object must be enclosed in double quotation marks.

--validate Lists all zone members that are not part of the current zone enforcement table.

If a zone member is specified by Domain and Port (D,P) in the Admin Domain membership list, all zone elements specified with WWNs associated with that D,P are considered for zone enforcement. If a device WWN is specified in the Admin Domain member list, the corresponding D,P (the device the specified WWN is associated with) is not considered for zone enforcement.

The following operands are optional:

- f**
Specifies that zone members that are not enforceable should be expunged in the transaction buffer. This pruning operation affects both the transaction buffer and the defined buffer. You cannot specify a mode or a zone object together with the -f option.
- mode**
Specifies the zone database location. This operand is optional. If no mode option is specified, the validated output of all the three buffers is displayed. Supported mode flag values include the following:
- 0**
Uses the zone database from the current transaction buffer.
- 1**
Uses the zone database stored in persistent storage.
- 2**
Uses the currently enforced zone database.
- "zone_object"**
Specifies a zone object. A zone object can be a zone member, a zone alias, or a zone.
- i**
Lists all zone members for a given pattern without case distinction.
- pattern**
A POSIX-style regular expression used to match zone members. This operand is optional. Patterns can contain:
- A question mark (?) to match any single character.
 - An asterisk (*) to match any string of characters.
 - A range of characters to match any character within the range: for example, [0-9] or [a-f].

2. Creating and managing TI Zones

Use these commands to create and manage TI Zones.

- create**
Creates a TI Zone with specified options and port list.
- add**
Adds port list members and the failover option to existing TI zones.
- remove**
Removes port list members from existing zones. Removal of the last member from an active TI zone generates a warning. If the last member of a TI zone is removed, the TI zone name is removed from the defined TI zone list.

The following operands are supported:

-t *objecttype*

Specifies the zone object type. This operand is supported only with the **--create** option. To create a TI zone, the value is **ti**.

-o *optlist*

Specifies list of options to control activation, deactivation, and failover mode. If this option is not specified the zone is created, by default, with failover enabled, and the zone will be activated. This operand is supported only with the **--create** and **--add** options.

Valid values for *optlist* include the following:

a

Activates the specified zone.

d

Deactivates the specified zone.

n

Disables failover mode. In non-failover mode, when the last interswitch link (ISL) of a TI Zone goes offline and there is an alternative ISL, the alternative ISL is not used and the switch generates a state change notification (SCN) or a registered state change notification (RSCN) to indicate that no ISL is available. When the ISL of the TI Zone comes online again, the route is set up again and the switch generates another SCN or RSCN. TI zones with no-failover option are not supported in logical fabrics. TI zones defined in the Base Fabric for logical fabric traffic need to allow failover.

f

Enables failover mode. In failover mode, when the last ISL of a TI Zone goes offline and there is an alternative ISL, the alternative ISL is used and the switch does not generate any SCN or RSCN messages. If the ISL of that TI Zone comes online again, traffic is rerouted immediately to the original ISL.

name

Specifies the name of the zone to be created, added or deleted

-p *portlist*

Specifies the lists of ports to be included, added or removed from a TI zone. The syntax for *portlist* is "D,I" (Domain, Index). The port list must be enclosed in double quotation marks. List members must be separated by semicolons, followed by a space. When you create TI zones over FCR, for a TI within an Edge fabric use "-1" in "I" (of "D,I") in to denote Front and Translate phantom in the E_Port list. When creating a TI zone within the Backbone fabric specify "Port WWN" in *portlist* to indicate devices talking across fabrics.

--activate *name*

Activates the specified TI zones. If more than one zone is specified, the list of zone names must be enclosed in double quotation marks; zone names must be separated by semicolons.

--deactivate *name*

Deactivates the specified TI zones. If more than one zone is specified, the list of zone names must be enclosed in double quotation marks; zone names must be separated by semicolons.

- delete *name***
Deletes TI zones from the defined TI zone lists. This command prompts for confirmation.
- show [*name*]**
Displays zone name, port lists, failover option, configured status and, the active status for the specified zones. The configured status reflects the latest change that may or may not have been activated. The active status reflects the state that is activated by **cfgEnable**. Without any specified name, this command displays all TI zones in the defined configuration. The following operands are optional:
- ic**
Displays all TI zone configurations for a given zone name without case distinction.
 - ascending**
Displays the TI zone members in ascending order.
- showTlerrors**
Analyzes real and potential routing problems with the activated TI zoning set and prints a report. This command must be executed in the local domain and analyzes only that domain.
- showTltrunkerrors**
Displays details of the trunk members present in the TI zone and those not present in the TI zone.

3. Creating and managing RD Zones

- rdcreate**
Creates a RD Zone for the specified members. The following operands are required:
- host_wwn***
Specifies the port world wide name of the host.
 - target_wwn***
Specifies the port world wide name of the target.
 - vi_wwn***
Specifies the port world wide name of the virtual initiator (VI).
 - vt_wwn***
Specifies the port world wide name of the virtual target (VT).
- policy restartable* | *nonrestartable***
Specifies the policy as either **restartable** or **nonrestartable**. A restartable policy causes traffic flow to revert to the physical host and target configuration in the event the virtual device should fail. When the policy is specified as nonrestricted and one of the virtual devices goes offline, the physical devices are considered offline and no traffic is allowed between the original host and target.
- FCR | noFCR**
Indicates whether or not this RD zone should function across a Fibre Channel router (FCR). Specify **FCR** to support FCR. Specify **noFCR** if you do not wish to support FCR.
- rddelete *name***
Deletes the RD Zone specified by ***name***. RD Zones are unique in that the zone name is not user-defined. It is derived based on members and the specified configuration. Use **cdfgShow** to display currently implemented RD zone by name.

Examples 1. Specific zone operation commands

The following example shows a scenario of an invalid configuration. If you attempt to create a zone z10 with an alias a10, then create a zone with name a10, z10 expects a10 to be an alias and results in an invalid configuration.

To validate all zones in the currently enforced zone database:

```
switch:admin> zonecreate z10,a10

switch:admin> zonecreate a10,1,2

switch:admin> zone --validate
Defined configuration:
zone: a10 1,2*
zone: z10 a10~

Effective configuration:
No Effective configuration: (No Access)
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

The reason for not being in the current enforcement table could be The device is not online.

To delete all references associated with zone member 100,5:

```
switch:admin> zone --expunge"100,5"
You are about to expunge one configuration
or member. This action could result in removing
many zoning configurations recursively.
[Removing the last member of a configuration
removes the configuration.]
Do you want to expunge the member?
(yes, y, no, n): [no] yes
```

To validate all zones in the zone database in the current transaction buffer:

```
switch:admin> zone --validate -m 0
Defined configuration:
cfg:   t_r_a_f_f_i_c_i_s_o_c_fg
       bluezone; greenzone
cfg:   ticonfig
       regzone
zone:  bluezone
       1,1*; 1,2*
zone:  greenzone 1,1*; 20:01:00:05:1e:01:23:e0*
zone:  regzone 1,4*; 1,5*
zone:  t_r_a_f_f_i_c_i_s_o_prop_zn
       1,3*; 2,3*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To validate all zones in the zone database in the persistent storage (defined configuration):

```
switch:admin> zone --validate -m 1
Defined configuration:
cfg:   t_r_a_f_f_i_c_i_s_o_c_fg
       bluezone; greenzone
cfg:   ticonfig
       regzone
zone:  bluezone
       1,1*; 1,2*
zone:  greenzone 1,1*; 20:01:00:05:1e:01:23:e0*
zone:  regzone 1,4*; 1,5*
zone:  t_r_a_f_f_i_c_i_s_o_prop_zn
       1,3*; 2,3*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To validate all zones in the zone database in the effective configuration:

```
switch:admin> zone --validate -m 2
Effective configuration:
cfg:   ticonfig
zone:  regzone 1,4*
       1,5*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To prune all the zone members that are not enforceable:

```
switch:admin> zone --validate -f
You are about to prune the zone configurations,
based on zone --validate output.
Do you want to prune the zone
configurations (yes, y, no, n): [no] y
```

To validate the zone members beginning with "gre", regardless of the case:

```
switch:admin> zone --validate -i gre*
Defined configuration:
zone:  GREEN 44,4; 21:00:00:20:37:0c:71:02; 8,9
zone:  green 2,2*; 2,3*; 21:00:00:20:37:0c:76:8c*

Effective configuration:
zone:  green 2,2*
       2,3*
       21:00:00:20:37:0c:76:8c*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

2. Traffic isolation zone commands

To create an activated traffic isolation zone with failover enabled (default).

```
switch:admin> zone --create -t ti purplezone \
-p "1,1; 2,4; 1,8; 2,6"
```

To create a deactivated traffic isolation zone with failover disabled:

```
switch:admin> zone --create -t ti -o dn purplezone \
-p "1,1; 2,4; 1,8; 2,6"
```

To add an E_Port and N_Port member as a *portlist* to an existing TI zone:

```
switch:admin> zone --add purplezone \
-p "3,4; 3,6"
```

To add the option to disable failover for a TI zone:

```
switch:admin> zone --add -o n purplezone
```

To add the option to enable failover for a TI zone:

```
switch:admin> zone --add -o f greenzone" -p "3,4"
```

To remove portlist members from an existing TI zone.

```
switch:admin> zone --remove bluezone -p "3,4; 3,6"
```

To activate the TI zone "bluezone":

```
switch:admin> zone --activate bluezone
```

To deactivate the TI zone "purplezone":

```
switch:admin> zone --deactivate purplezone
```

To delete the TI zone "bluezone":

```
switch:admin> zone --delete bluezone
```

To display all TI zones in the defined configuration:

```
switch:admin> zone --show
Defined TI zone configuration:

TI Zone Name:    ti_blue
Port List:  4,55; 10:00:00:00:00:01:00:01; \
            10:00:00:00:00:01:00:04; 3,9

Configured Status: Activated / Failover-Disabled
Enabled Status: Deactivated

TI Zone Name:    ti_red
Port List:  4,4; 5,5; 3,3

Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
```

To display all TI zones in the defined configuration in ascending order:

```
switch:admin> zone --show -ascending

Defined TI zone configuration:

TI Zone Name:   ti_blue
Port List:      10:00:00:00:00:01:00:01; \
                10:00:00:00:00:01:00:04; 3,9; 4,55

Configured Status: Activated / Failover-Disabled
Enabled Status: Deactivated

TI Zone Name:   ti_red
Port List:      3,3; 4,4; 5,5

Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
```

To display a specified zone in the defined configuration:

```
switch:admin> zone --show ti_red

Defined TI zone configuration:

TI Zone Name:   ti_red
Port List:      4,4; 5,5; 3,3

Configured Status: Activated / Failover-Disabled
Enabled Status: Deactivated
```

To display members for the zone "ti_red" in ascending order:

```
switch:admin> zone --show ti_red -ascending

Defined TI zone configuration:

TI Zone Name:   ti_red
Port List:      3,3; 4,4; 5,5

Configured Status: Activated / Failover-Disabled
Enabled Status: Deactivated
```

To display members for the zone "TI_zone", regardless of the case:

```
switch:admin> zone --show -ic TI_zone*

Defined TI zone configuration:

TI Zone Name:   TI_zone
Port List:      7,8

Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated

TI Zone Name:   ti_zone
Port List:      3,3

Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated
```

To create an FCR TI zone within an Edge fabric where a host should talk to target1 and target4 through port number 2 on an Edge fabric switch with a domain ID of 3. (in the example, "3,1" is the host, and the remaining members are E_Ports):

```
switch:admin> zone --create -t ti fcr_edge_ti_zone \
-p "3,1; 3,2; 4,-1; 5, -1"
```

To create an FCR TI zone within a Backbone where a host, target1, and target4 communicate over VE_Ports consisting of FCR1 port number 4 and FCR2 port number 7:

```
switch:admin> zone --create -t ti fcr_ti_zone \
-p "1,1; 2,1; host_PWWN; target1_PWWN; \
target4_PWWN; 1,4; 2,7"
```

To troubleshoot TI zone problems:

```
switch:admin> zone --showTIerrors
My Domain: 1

Error type:          ERROR
Affected Remote Domain: 2
Affected Local Port: 3
Affected TI Zones:   ti1, ti2
Affected Remote Ports: 6, 7, 8, 9
```

To display details of the trunk members present in the TI zone and those not present in the TI zone:

```
switch:admin> zone --showTItrunkerrors
TI Zone Name: brackets

E-Port Trunks
Trunk members in TI zone: 16 18
Trunk members not in TI zone: 17

F-Port Trunks
Trunk members in TI zone: 4 5
Trunk members not in TI zone: 6

TI Zone Name: loop

E-Port Trunks
Trunk members in TI zone: 0
Trunk members not in TI zone: 1

TI Zone Name: operand

E-Port Trunks
Trunk members in TI zone: 8
Trunk members not in TI zone: 9 10

E-Port Trunks
Trunk members in TI zone: 16
Trunk members not in TI zone: 17 18
```

3. Frame redirect zone commands

To create an RD Zone, given a host (10:10:10:10:10:10:10:10), target (20:20:20:20:20:20:20:20), VI (30:30:30:30:30:30:30:30), and VT (40:40:40:40:40:40:40:40):

```
switch:admin> zone --rdcreate \
    10:10:10:10:10:10:10:10 20:20:20:20:20:20:20:20 \
    30:30:30:30:30:30:30:30 40:40:40:40:40:40:40:40 \
    restartable noFCR
```

This command creates the following zone objects:

- RD zone "red_0917_10_10_10_10_10_10_10_10_10_10_10_10_20_20_20_20_20_20_20_20_20_20_20_20", with a restricted policy and no FCR support.
- The base zone object, "red_____base".
- The RD zone configuration, "r_e_d_i_r_c__fg".

To display the newly created zone objects:

```
switch:admin> cfgshow
Defined configuration:
cfg:   myHTcfg myHostTarget
cfg:   r_e_d_i_r_c__fg
       red_____base;
red_0917_00_3f_3f_3f_23_24_25_26_3f_3f_3f_30_32_00_00_00
zone:  myHostTarget
      00:3f:3f:3f:23:24:25:26; 3f:3f:3f:30:32:00:00:00
zone:red_0917_00_3f_3f_3f_23_24_25_26_3f_3f_3f_30_32_00_00_00
      00:3f:3f:3f:23:24:25:26; 3f:3f:3f:30:32:00:00:00;
      3f:3f:3f:30:30:00:00:00; 3f:3f:3f:30:31:00:00:00
zone:  red_____base
      00:00:00:00:00:00:00:01; 00:00:00:00:00:00:00:02;
      00:00:00:00:00:00:00:03; 00:00:00:00:00:00:00:04

Effective configuration:
cfg:   myHTcfg
zone:  myHostTarget
      00:3f:3f:3f:23:24:25:26
      3f:3f:3f:30:32:00:00:00
```

To delete an RD Zone named "red_0917_10_10_10_10_10_10_10_10_10_10_10_10_20_20_20_20_20_20_20_20_20_20_20_20":

```
switch:admin> zone --rddelete \
    red_0917_10_10_10_10_10_10_10_10_10_10_10_10_20_20_20_20_20_20_20_20_20_20_20_20
```

SEE ALSO [ad, zoneHelp](#)

zoneAdd

Adds a member to the zone or Peer Zone.

Synopsis `zoneadd "zoneName", "member[;member...]"`

`zoneadd --peerzone "zoneName" -principal "principal[;principal...]"
-members "member[;member...]"`

Description Use this command to add one or more members to an existing zone or Peer Zone.

This command changes the defined configuration. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the **cfgSave** command. For the change to become effective, enable the configuration with the **cfgEnable** command.

Target Driven Peer Zones are not allowed to be edited using this command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"zoneName"

Specify the name of an existing zone. Double quotation marks are optional.

"member"

Specify a member or list of members to be added. The list must be enclosed in double quotation marks. Members must be separated by semicolons. Members can be specified in one or more of the following ways:

- A switch domain and port index number pair; for example, "2, 20". Use **switchShow** for a listing of valid port index numbers.
- Node or port WWN.
- Zone alias name.
- Zone alias pattern. Use a pattern preceded by the wildcard symbol asterisk (*) to include members in the zone that matches the pattern.

--peerzone "zoneName"

Specifies the Peer Zone name to which one or more principal or members should be added.

-principal "principal"

Specify a principal or list of principal members to be added. The list must be enclosed in double quotation marks. Principal members must be separated by semicolons. You can specify either a WWN or D, I; zone alias is not a valid principal member.

-members "member"

Specify a member or list of members to be added to the Peer Zone. You can specify either a WWN or D, I; zone alias is not a valid member.

Examples To add aliases for three disk arrays to "Blue_zone":

```
switch:admin> zoneadd "Blue_Zone", "array3; array4; array5"
```

To add aliases for all the arrays that matches a pattern:

```
switch:admin> zoneadd matt, "ze*;bond*;j*"
switch:admin> cfgshow
Defined configuration:
zone: matt 30:06:00:07:1e:a2:10:20; 3,2; zeus; \
bond; jake; jeff; jones
zone: sloth bawn; bolt; bond; brain; brit; bru; \
10:00:00:00:01:1e:20:20
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: brit 12,1
alias: bru 5,3; 12,4
alias: geek 2,7; 4,11; 20:10:00:05:1e:a9:20:87
alias: jake 4,7; 8,9; 14,11
alias: jeff 30:00:00:05:1e:a1:cd:02; \
40:00:00:05:1e:a1:cd:04
alias: jones 7,3; 4,5
alias: zeus 4,7; 6,8; 9,2

Effective configuration:
No Effective configuration: (No Access)
```

To add a member to a Peer Zone:

```
switch:admin> zoneadd --peerzone peerzone_wwn_mbrs \
-members "10:00:05:1e:a9:20:00:02"
```

SEE ALSO [zoneCreate](#), [zoneDelete](#), [zoneRemove](#), [zoneShow](#)

zoneCreate

Creates a zone or a Peer Zone.

Synopsis `zonecreate "zonename", "member[:member...]"`

`zonecreate --peerzone "zonename" -principal "principal[:principal...]"
[-members "member[:member...]"`

Description Use this command to create a new zone or a Peer Zone, or to create a "broadcast" zone.

A broadcast zone is a special zone that specifies the nodes that can receive broadcast traffic. This zone must be named "broadcast". Only one "broadcast" zone can exist within a fabric. This type of zone is enforced by the hardware; the switch controls the data transfer to a port.

This command changes the defined configuration. For the change to be preserved across switch reboots, save it to nonvolatile memory with the **cfgSave** command. For the change to become effective, enable the zone configuration with the **cfgEnable** command.

Target Driven Peer Zones are not allowed to be created using this command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"zonename"

Specify a unique name for the zone to be created. Double quotation marks are optional. A zone name must begin with a letter and followed by any number of letters, numbers, or underscore characters. Names are case-sensitive. For example, "Zone_1" and "zone_1" indicate different zones. Zone names are limited to 64 characters. Spaces are ignored.

"member"

Specify a member or a list of members to be included in the zone. The list must be enclosed in double quotation marks. Members must be separated by semicolons. The list must have at least one member; empty lists are not allowed.

A member can be specified in one or more of the following ways:

- **Domain and port index pair:** Specify a port by domain and port index, for example, "2, 20" specifies port index 20 on switch domain 2. When a zone member is specified by port index, all devices connected to that port are in the zone. If this port is an arbitrated loop, then all devices on the loop are in the zone. Use **switchShow** for a list of valid port index numbers.
- **World wide name:** Specify a world wide name as eight hex numbers separated by colons, for example "10:00:00:60:69:00:00:8a". Zoning compares the WWN with the node and port names presented by a device in a login frame (FLOGI or PLOGI). When a zone member is specified by node name, then all ports on that device are in the zone. When a zone member is specified by port name, then only that single device port is in the zone.
- **Zone alias name:** Specify a zone alias name using the same format as a zone name. Refer to **aliCreate** command help for zone alias naming requirements.
- **Zone alias pattern:** Use a pattern preceded by the wildcard symbol asterisk (*) to include members in the zone that matches the pattern.

--peerzone "zoneName"

Creates a new Peer Zone.

-principal "principal"

Specify a principal or list of principal members to be included in the Peer Zone. The list must be enclosed in double quotation marks. Principal members must be separated by semicolons. You can specify either a WWN or D, I; zone alias is not a valid principal member.

-members "member"

Specify a member or list of members to be included in the Peer Zone. You can specify either a WWN or D, I; zone alias is not a valid member. This operand is optional.

When creating a zone, you can combine different ways of specifying zone members. For example, a zone defined with the following members: "2,12; 2,14; 10:00:00:60:69:00:00:8a" contains all devices connected to switch 2, ports 12 and 14, and to the device with the world wide name "10:00:00:60:69:00:00:8a" (either node name or port name), at the port in the fabric to which it is connected.

Examples

To create three zones using a combination of port numbers and zone aliases:

```
switch:admin> zonecreate "Purple_zone", "1,0"

switch:admin> zonecreate "Blue_zone", "1,1; array1; 1,2; array2"

switch:admin> zonecreate "Green_zone", "1,0; 1,2; array2"
```

To create zone with the specified zone alias and to include the members that matches a pattern:

```
switch:admin> zonecreate sloth, "b*; 10:00:00:00:01:1e:20:20"
switch:admin> cfgshow
Defined configuration:
zone: matt 30:06:00:07:1e:a2:10:20; 3,2
zone: sloth bawn; bolt; bond; brain; brit; \
      bru; 10:00:00:00:01:1e:20:20
alias: bawn 3,5; 4,8
alias: bolt 10:00:00:02:1f:02:00:01
alias: bond 10:00:05:1e:a9:20:00:01; 3,5
alias: brain 11,4; 22,1; 33,6
alias: brit 12,1
```

```

alias:      bru      5,3; 12,4
alias:      geek     2,7; 4,11; 20:10:00:05:1e:a9:20:87
alias:      jake     4,7; 8,9; 14,11
alias:      jeff     30:00:00:05:1e:a1:cd:02; \
                    40:00:00:05:1e:a1:cd:04
alias:      jones    7,3; 4,5
alias:      zeus     4,7; 6,8; 9,2

```

To create Peer Zones:

```

switch:admin> zonecreate --peerzone peerzone_wwn_mbrs \
                    -principal "10:00:00:00:01:1e:20:20" -members "10:00:00:02:1f:02:00:01;10:00:05:1e:a9:20:00:01"

switch:admin> zonecreate --peerzone peerzone_di_mbrs -principal "10,1" -members "20,1;20,2"

```

SEE ALSO [zoneAdd](#), [zoneDelete](#), [zoneRemove](#), [zoneShow](#)

zoneDelete

Deletes a zone or Peer Zone.

Synopsis `zonedeleter "zonename"`

Description Use this command to delete a zone or Peer Zone.

This command changes the defined configuration. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory using the **cfgSave** command. For the change to become effective, enable the configuration with the **cfgEnable** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operand is required:

`"zonename"`

Specify the name of the zone to be deleted. Quotation marks are optional.

Examples To delete the zone "Blue_zone":

```
switch:admin> zonedeleter "Blue_zone"
```

To delete the Peer Zone "peerzone_di_mbrs":

```
switch:admin> zonedeleter "peerzone_di_mbrs"
```

SEE ALSO [zoneAdd](#), [zoneCreate](#), [zoneRemove](#), [zoneShow](#)

zoneHelp

Displays a description of zoning commands.

Synopsis `zonehelp`

Description Use this command to display short descriptions of zoning commands.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands None

Examples To display zone command help information:

```
switch:admin> zonehelp
aliadd          Add a member to a zone alias
alcreate        Create a zone alias
aldelete        Delete a zone alias
alremove        Remove a member from a zone alias
alshow          Print zone alias information
bootluncfg      Configure boot LUN for an HBA
cfgactvshow     Display Effective zone configuration \
                information
cfgadd          Add a member to a configuration
cfgclear        Clear all zone configurations
cfgcreate       Create a zone configuration
cfgdelete       Delete a zone configuration
cfgdisable      Disable a zone configuration
cfgenable       Enable a zone configuration
cfgremove       Remove a member from a configuration
cfgsave         Save zone configurations in flash
cfgshow         Print zone configuration information
cfgsize         Print size details of zone database
cfgtransabort   Abort zone configuration transaction
cfgtransshow    Print zone configurations in transaction buffer
defzone         Activates or deactivates a default zone
                configuration.
msfr            Create a MSFR Zone (internal use only)
nszonemember    Display the information of all the online devices
                which are zoned with the given device.
openfr          Create a MSFR Zone (internal use only)
zone            Copies/Removes/Validates zone objects
zoneadd         Add a member to a zone
zonecreate      Create a zone
zonedel         Delete a zone
zonehelp        Print zoning help info
zoneobjectcopy  Copies a zone object
zoneobjectexpunge Expunges a zone object
zoneobjectrename Rename a zoning Object
zoneremove      Remove a member from a zone
zoneshow        Print zone information
```

SEE ALSO None

zoneObjectCopy

Copies a zone object.

Synopsis `zoneObjectCopy "objectName", "newName"`

Description Use this command to make a copy of an existing zone object and give it a new name. The resulting object is of the same type as the original object. You can use this command for all zone object types, including `cfg`, `zone`, and `alias`.

Target Driven Peer Zones are not allowed to be edited using this command.

This command changes the defined configuration. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the `cfgSave` command. For the change to become effective, enable the configuration with the `cfgEnable` command.

This command is not applicable for special zones, such as TI zones or RD zones.

Notes When FCS policy is enabled, this command can be issued only from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operands are required:

"objectName"

The name of the object that you want to copy. Quotation marks are optional.

"newName"

The name of the object that you want created. Quotation marks are optional.

A zone configuration name must begin with a letter followed by any number of letters, numbers, and underscores. Names are case-sensitive. For example, "Cfg_1" and "cfg_1" are different zone configurations. Blank spaces are ignored.

Refer to the `zoneCreate` command for more information on name and member specifications

Examples To create a configuration containing three zones:

```
switch:admin> cfgshow ""
  cfg:  USA_cfg  Red_zone; White_zone; Blue_zone

switch:admin> zoneobjectcopy "USA_cfg", "UK_cfg"

switch:admin> cfgshow ""
  cfg:  UK_cfg  Red_zone; White_zone; Blue_zone
  cfg:  USA_cfg  Red_zone; White_zone; Blue_zone
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#), [zoneObjectRename](#)

zoneObjectExpunge

Expunges a zone object.

Synopsis `zoneObjectExpunge "objectName"`

Description Use this command to expunge a zone object. In addition to deleting the object, this command also removes the object from the member lists of all other objects. After successful execution of this command, the specified object no longer exists in the database. You can use this command for all zone object types, including `cfg`, `zone`, and `alias`.

Target Driven Peer Zones are not allowed to be edited using this command.

This command changes the defined configuration. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the `cfgSave` command. For the change to become effective, enable the configuration with the `cfgEnable` command.

Notes When FCS policy is enabled, this command can be issued only from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands This command has the following operand:

`"objectName"`

The name of the object that you want to expunge. Quotation marks are optional. This operand is required.

Examples To create a configuration containing three zones:

```
switch:admin> cfgshow
Defined configuration:
cfg:   USA_cfg Red_zone; White_zone; Blue_zone
zone:  Blue_zone
      1,1; array1; 1,2; array2
zone:  Red_zone
      1,0; loop1
zone:  White_zone
      1,3; 1,4
alias: array1  21:00:00:20:37:0c:76:8c; \
      21:00:00:20:37:0c:71:02
alias: array2  21:00:00:20:37:0c:76:22; \
      21:00:00:20:37:0c:76:28
alias: loop1   21:00:00:20:37:0c:76:85; \
      21:00:00:20:37:0c:71:df
```

```
switch:admin> zoneobjectexpunge "White_zone"
```

```
switch:admin> cfgshow
Defined configuration:
cfg:   USA_cfg   Red_zone; Blue_zone
zone:  Blue_zone
      1,1; array1; 1,2; array2
zone:  Red_zone
      1,0; loop1
```

zoneObjectExpunge

```
alias: array1 21:00:00:20:37:0c:76:8c; \  
21:00:00:20:37:0c:71:02  
alias: array2 21:00:00:20:37:0c:76:22; \  
21:00:00:20:37:0c:76:28  
alias: loop1 21:00:00:20:37:0c:76:85; \  
21:00:00:20:37:0c:71:df
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#), [zoneObjectCopy](#), [zoneObjectRename](#)

zoneObjectRename

Renames a zone object.

Synopsis `zoneObjectRename "objectName", "newName"`

Description Use this command to rename a zone object. You can use this command for all zone object types, including `cfg`, `zone`, and `alias`.

Target Driven Peer Zones are not allowed to be edited using this command.

This command changes the defined configuration. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the `cfgSave` command. For the change to become effective, enable the configuration must be enabled with the `cfgEnable` command.

This command is not applicable for special zones, such as TI zones or RD zones.

Notes When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

Operands The following operands are required:

"objectName"

The name of the object you want to rename.

"newName"

The new name of the object.

A zone configuration name must begin with a letter that can be followed by any number of letters, numbers, and underscores. Names are case-sensitive. For example, "Cfg_1" and "cfg_1" are different zone configurations. Quotation marks are optional. Spaces are ignored.

Refer to the `zoneCreate` command for more information on name and member specifications.

Examples To create a configuration containing three zones:

```
switch:admin> cfgshow ""
  cfg:  USA_cfg  Red_zone; White_zone; Blue_zone

switch:admin> zoneobjectrename "USA_cfg", "UK_cfg"

switch:admin> cfgshow ""
  cfg:  UK_cfg  Red_zone; White_zone; Blue_zone
```

SEE ALSO [cfgAdd](#), [cfgClear](#), [cfgDelete](#), [cfgDisable](#), [cfgEnable](#), [cfgRemove](#), [cfgSave](#), [cfgShow](#), [zoneObjectCopy](#)

zoneObjectReplace

Replaces zone members.

Synopsis `zoneobjectreplace oldmember newmember`

`zoneobjectreplace --help`

Description Use this command to replace the existing member of a zone or a Peer Zone with a new member. This command can be used to replace members of an alias, but an alias itself cannot be replaced. This command is not applicable for Traffic Isolation (TI) Zones, and Frame Redirect (RD) Zones.

The property member of a Peer Zone cannot be replaced using this command.

Target Driven Peer Zones are not allowed to be edited using this command.

For the change to become effective, enable the configuration with the **cfgEnable** command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

oldmember

Specifies the the zone member that has to be replaced. You can specify WWN or Domain,Index (D,I).

newmember

Specifies the new zone member with which to replace the old zone member. You can specify WWN or D,I.

--help

Displays the command usage.

Examples To replace zone member:

```
switch:admin> zoneobjectreplace \  
    "20:00:00:05:1e:a1:af:b2" "10:00:00:05:1e:a1:10:c1"  
switch:admin> cfgshow  
Defined configuration:  
cfg:   cfg       lsan_m  
cfg:   cfg1      zone2; zone3  
zone:  lsan_m    20:15:00:05:1e:a2:f9:b1; 20:13:00:05:1e:a2:f9:b1  
zone:  zone1     2,3; 20:00:00:05:1e:a1:af:b1; zali_1  
zone:  zone2     5,7; 2,9  
zone:  zone3     10:00:00:05:1e:a1:10:c1; 20:10:00:05:1e:a9:20:12
```

```
alias: zali_1 10:00:00:05:1e:a1:ef:b9; 2,3; 1,11
```

Effective configuration:

```
cfg:  cfg
```

```
zone:  lsan_m 20:15:00:05:1e:a2:f9:b1 20:13:00:05:1e:a2:f9:b1
```

SEE ALSO [zoneAdd](#), [zoneCreate](#), [zoneDelete](#), [zoneRemove](#), [zoneShow](#)

zoneRemove

Removes a member and principal (Peer Zone) from a zone.

Synopsis `zoneremove "zonename", "member[:member...]"`

`zoneremove --peerzone "zonename" -principal "principal[:principal...]"`
`-members "member[:member...]"`

Description Use this command to remove one or more members from an existing zone, and to remove one or more members or principal members from a Peer Zone.

If all members are removed, the zone is deleted.

If all members (principal and non-principal members) are removed, the Peer Zone is deleted. If all non-principal members are removed but there is still a principal member, the Peer Zone will still exist.

This command changes the defined configuration. For the change to be preserved across switch reboots, save the configuration to nonvolatile memory with the **cfgSave** command. For the change to become effective, enable the configuration with the **cfgEnable** command.

Target Driven Peer Zones are not allowed to be edited using this command.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When an FCS policy is enabled, this command can be issued only from the primary FCS switch.

Operands The following operands are required:

"zonename"

Specify the name of the zone from which to remove a member or principal (Peer Zone only). Double quotation marks are optional.

"member"

Specify a member or a list of members to be removed from the specified zone. The list must be enclosed in double quotation marks. Members must be separated by semicolons. A member can be one or more of the following:

- A switch domain and port index pair: for example, "2,20". Use **switchShow** for a list of valid port index numbers.
- A world wide name
- A zone alias name
- Zone alias pattern. Use a pattern preceded by the wildcard symbol asterisk (*) to include members in the zone that matches the pattern.

--peerzone "zonename"

Specifies the Peer Zone name from which one or more principal or members are to be removed.

-principal "principal"

Specify a principal or list of principal members to be removed from a Peer Zone. The list must be enclosed in double quotation marks. Principal members must be separated by semicolons. You can specify either a WWN or D, I; zone alias is not a valid principal member.

-members "member"

Specify a member or list of members to be removed from a Peer Zone. You can specify either a WWN or D, I; zone alias is not a valid member.

Examples

To remove "array2" from "Blue_zone":

```
switch:admin> zoneremove "Blue_zone", "array2"
```

To remove all the members and aliases matching a pattern:

```
switch:admin> zoneremove matt,"30:06:00:07:1e:a2:10:20; ja*; 3,2"
switch:admin> cfgshow
Defined configuration:
zone: matt   zeus; bond; jeff; jones
zone: sloth  bawn; bolt; bond; brain; brit; \
            bru; 10:00:00:00:01:1e:20:20
alias:      bawn   3,5; 4,8
alias:      bolt   10:00:00:02:1f:02:00:01
alias:      bond   10:00:05:1e:a9:20:00:01; 3,5
alias:      brain  11,4; 22,1; 33,6
alias:      brit   12,1
alias:      bru    5,3; 12,4
alias:      geek   2,7; 4,11; 20:10:00:05:1e:a9:20:87
alias:      jake   4,7; 8,9; 14,11
alias:      jeff   30:00:00:05:1e:a1:cd:02; \
            40:00:00:05:1e:a1:cd:04
alias:      jones  7,3; 4,5
alias:      zeus   4,7; 6,8; 9,2

Effective configuration:
No Effective configuration: (No Access)
```

To remove a member from a Peer Zone:

```
switch:admin> zoneremove --peerzone peerzone_wwn_mbrs \
-members "10:00:05:1e:a9:20:00:02"
```

SEE ALSO [zoneAdd](#), [zoneCreate](#), [zoneDelete](#), [zoneShow](#)

zoneShow

Displays zone information.

Synopsis

```
zonestow [--sort] [pattern][, mode]
zonestow [--ic] [pattern][, mode]
zonestow [--alias] [-ic] [pattern]
zonestow [--transdiffs]
zonestow [--transdiffonly]
zonestow [--validate [pattern] [,mode]]
zonestow [--peerzone [all | user | target]]
zonestow --help
```

Description Use this command to display zone configuration information. This command includes sorting and search options to customize the output. If a pattern is specified, the command displays only matching match zone configuration names in the defined configuration. When used without operands, the command displays all zone configuration information for the Defined and the Effective configuration.

Use the **--transdiffs** and **--transdiffonly** options to view changes in the current transaction.

The command output displays the changes in the current transaction by the following notations:

- An asterisk(*) before any tag indicates a change in that zone, zone configuration, alias or any other entity in the zone configuration.
- A plus(+) before any entity indicates that it is a newly added entity.
- A minus(-) before any entity indicates that it is a deleted entity.

Refer to **cfgShow** for a description of this display.

Notes The execution of this command is subject to Virtual Fabric restrictions that may be in place. Refer to "Using Fabric OS Commands" and "Command Availability" for details.

When FCS policy is enabled, this command can be issued on any FCS switch in the fabric.

Operands This command has the following operands:

--sort

Displays D,I zone members and WWNs in ascending order. For Peer Zones, the principal and peer members are sorted separately.

--ic

Displays all zone configuration names for a given pattern without case distinction.

pattern

A POSIX-style regular expression used to match zone configuration names. This operand is optional. Patterns can contain:

- A question mark (?) to match any single character.
- An asterisk (*) to match any string of characters.
- A range of characters to match any character within the range: for example, [0-9] or [a-f].

mode

Specify 0 to display the contents of the transaction buffer (the contents of the current transaction), or specify 1 to display the contents of the nonvolatile memory. The default value is 0. This operand is optional.

--alias

Displays all the zones containing the given alias pattern. The *pattern* operand is mandatory. The pattern can contain:

- A question mark (?) to match any single character.
- An asterisk (*) to match any string of characters.
- A range of characters to match any character within the range: for example, [0-9] or [a-f].

--transdiffs

Displays the changes in the current transaction.

--transdiffsonly

Displays only the changes in the current transaction.

--validate

Displays the validated output for the specified zone configuration or for all the zone configurations in the zone database.

zone_name

Specify the exact zone name or a portion of the zone name followed by an asterisk (*) at the end to match any string of characters.

mode

Specifies the zone database location. This operand is optional. If no mode option is specified, the validated output of all the three buffers is displayed. Supported mode flag values include the following:

0 | trans

Uses the zone database from the current transaction buffer.

1 | def

Uses the zone database stored in persistent storage.

2 | eff

Uses the currently enforced zone database.

--peerzone

Displays the Peer Zone configuration information (Peer Zone name, Peer Zone type, property member, principal members, and peer members) for the Defined and the Effective configuration. The following operands are valid:

all

Displays configuration information for all Peer Zones.

user

Displays configuration information for the user created Peer Zones.

target

Displays configuration information for the target created Peer Zones.

--help

Displays the command usage.

Examples

To display all zones:

```
switch:admin> zonestow
Defined configuration:
  cfg:  cfg1  red
  zone:  blue  44,5; 10:00:00:00:00:01:00:00; 3,4
  zone:  red   3,4; 1,2; 4,5; 2,3

Effective configuration:
  cfg:  cfg1
  zone:  red   3,4
           1,2
           4,5
           2,3
```

To sort the zones in ascending order:

```
switch:admin> zoneShow --sort
Defined configuration:
  cfg:  cfg1  red
  zone:  blue  10:00:00:00:00:01:00:00; 3,4; 44,5
  zone:  red   1,2; 2,3; 3,4; 4,5

Effective configuration:
  cfg:  cfg1
  zone:  red   1,2
           2,3
           3,4
           4,5
```

To display the red zone only using pattern search:

```
switch:admin> zonestow red
zone:  red   3,4; 1,2; 4,5; 2,3
```

To combine a pattern search with the sorting option:

```
switch:admin> zonestow --sort red
zone:  red   1,2; 2,3; 3,4; 4,5
```

To display the filtered content of the transaction buffer:

```
switch:admin> zonestow red,1
zone:  red   3,4; 1,2; 4,5; 2,3
```

To display the filtered and sorted content of the transaction buffer:

```
switch:admin> zonestow --sort red,1
zone:  red   1,2; 2,3; 3,4; 4,5
```


To display all green zones using pattern search, regardless of the case:

```
switch:admin> zonestow --ic GREEN*
zone: GREEN 44,4; 21:00:00:20:37:0c:71:02; 8,9
zone: green 2,2; 2,3; 21:00:00:20:37:0c:76:8c
```

To display the zone members of alias "ali10":

```
switch:admin> zonestow --alias ali10
zone: zone20 ali10
```

To display the zone members of aliases beginning with "ali1":

```
switch:admin> zonestow --alias ali1*
zone: zone1 1,2; ali12; 3,3; 30:04:00:05:1e:0b:55:0f; 30:05:00:05:1e:0b:55:0f
zone: zone 15 9,8; 2,5; 3,3; ali13
zone: zone 17 ali1
zone: zone 20 ali10
```

To display the changes in the current transaction:

```
switch:admin> zonestow --transdiffs
Defined configuration:
cfg: fabric_cfg
    green_zone
zone: blue_zone
    21:00:00:20:37:0c:76:8c
*zone: green_zone
    21:00:00:20:37:0c:76:01; -1,4
*zone: red_zone
    3,3; 3,4; +5,5; +5,6

Effective configuration:
cfg: fabric_cfg
zone: green_zone
    21:00:00:20:37:0c:76:01
    1,4
```

To display only the changes in the current transaction:

```
switch:admin> zonestow --transdiffonly
*zone: green_zone
    21:00:00:20:37:0c:76:01; -1,4
*zone: red_zone
    3,3; 3,4; +5,5; +5,6
```

To display validated output of all zones in the zone database:

```
switch:admin> zonestow --validate
Defined configuration:
cfg: cfg1 zone1; zone10; zone2
zone: zone1 20:1c:00:05:1e:57:b1:c6*; 20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*

Effective configuration:
cfg: cfg1
zone: zone1 20:1c:00:05:1e:57:b1:c6*
```

zoneShow

```
20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6
      20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6
      20:1f:00:05:1e:57:b1:c6*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To display validated output for a zone:

```
switch:admin> zoneshow --validate zone1
Defined configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*; 20:1d:00:05:1e:57:b1:c6

Effective configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*
      20:1d:00:05:1e:57:b1:c6
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To display validated output for zone members beginning with "zone":

```
switch:admin> zoneshow --validate zone*
Defined configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*; 20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
zone: zone200 20:1d:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*

Effective configuration:
zone: zone1 20:1c:00:05:1e:57:b1:c6*
      20:1d:00:05:1e:57:b1:c6
zone: zone10 20:1e:00:05:1e:57:b1:c6
      20:1f:00:05:1e:57:b1:c6*
zone: zone2 20:03:00:05:1e:57:b1:c6
      20:1f:00:05:1e:57:b1:c6*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To display validated output for a zone in the current transaction buffer:

```
switch:admin> zoneshow --validate zone200 O
Defined configuration:
zone: zone200 20:1d:00:05:1e:57:b1:c6; 20:1f:00:05:1e:57:b1:c6*
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

To display the complete Peer Zoning configuration:

```
switch:admin> zoneshow --peerzone all
Defined configuration:
zone: peerzone_di_mbrs
  Property Member: 00:02:00:00:00:02:00:01
  Created by: User
  Principal Member(s):
    10,1
  Peer Member(s):
    20,1; 20,2
zone: peerzone_wwn_mbrs
  Property Member: 00:02:00:00:00:03:00:01
  Created by: User
  Principal Member(s):
    10:00:00:00:01:1e:20:20
  Peer Member(s):
    10:00:00:02:1f:02:00:01; \
    10:00:05:1e:a9:20:00:01
Effective configuration:
zone: peerzone_wwn_mbrs
  Property Member: 00:02:00:00:00:03:00:01
  Created by: User
  Principal Member(s):
    10:00:00:00:01:1e:20:20
  Peer Member(s):
    10:00:00:02:1f:02:00:01
    10:00:05:1e:a9:20:00:01
1 Peer Zones in Eff Cfg
```

SEE ALSO [zoneAdd](#), [zoneCreate](#), [zoneDelete](#), [zoneRemove](#)

zoneShow

Primary FCS Commands

Primary FCS commands

Table 4 summarizes the commands that are available only on the primary Fabric Configuration Server (FCS) when FCS policy is enabled.

TABLE 4 Primary FCS commands

| Command | Description |
|---------------------------|---|
| aliAdd | Must be run from the primary FCS switch. |
| aliCreate | Must be run from the primary FCS switch. |
| aliDelete | Must be run from the primary FCS switch. |
| aliRemove | Must be run from the primary FCS switch. |
| aliShow | Can be run on all FCS switches. |
| cfgAdd | Must be run from the primary FCS switch. |
| cfgClear | Must be run from the primary FCS switch. |
| cfgCreate | Must be run from the primary FCS switch. |
| cfgDelete | Must be run from the primary FCS switch. |
| cfgDisable | Must be run from the primary FCS switch. |
| cfgEnable | Must be run from the primary FCS switch. |
| cfgRemove | Must be run from the primary FCS switch. |
| cfgSave | Must be run from the primary FCS switch. |
| cfgShow | Can be run on all FCS switches. |
| cfgTransAbort | Must be run from the primary FCS switch. |
| cfgTransShow | Must be run from the primary FCS switch. |
| date | This command can be run on all switches to view the current date. You can only modify the date from the primary FCS switch. |
| defZone | defzone --show can be run on all switches. All other options must be run from the primary FCS switch. |
| msPIClearDB | Must be run from the primary FCS switch. |
| msPIMgmtActivate | Must be run from the primary FCS switch. |
| msPIMgmtDeactivate | Must be run from the primary FCS switch. |
| msTdDisable | msTdDisable "ALL" must be run from the primary FCS switch. |
| msTdEnable | msTdEnable "ALL" must be run from the primary FCS switch. |
| secPolicyAbort | Must be run from the primary FCS switch. |
| secPolicyActivate | Must be run from the primary FCS switch. |
| secPolicyAdd | Must be run from the primary FCS switch. |
| secPolicyCreate | Must be run from the primary FCS switch. |
| secPolicyDelete | Must be run from the primary FCS switch. |
| secPolicyDump | Can be run on all FCS switches. |

TABLE 4 Primary FCS commands (Continued)

| Command | Description |
|-------------------------|--|
| secPolicyFCSMove | Must be run from the primary FCS switch. |
| secPolicyRemove | Must be run from the primary FCS switch. |
| secPolicySave | Must be run from the primary FCS switch. |
| secPolicyShow | Can be run on all FCS switches. |
| snmpConfig | Can be run on all FCS switches. |
| tsClockServer | Can be run on all switches to view the NTP server's IP address. You can only modify the NTP server's IP address on the primary FCS switch. |
| zoneAdd | Must be run from the primary FCS switch. |
| zoneCreate | Must be run from the primary FCS switch. |
| zoneDelete | Must be run from the primary FCS switch. |
| zoneObjectCopy | Must be run from the primary FCS switch. |
| zoneObjectExpung | Must be run from the primary FCS switch. |
| zoneObjectRename | Must be run from the primary FCS switch. |
| zoneRemove | Must be run from the primary FCS switch. |

Command Availability

- [Command validation checks](#) 1183
- [Command Virtual Fabrics and Remote Domain execution restrictions](#) 1184

Command validation checks

Before a command is executed, it is validated against the following checks.

1. Active or Standby availability: On enterprise-class platforms systems, checks that the command is available on the Control Processor (CP).
2. Role-Based Access Control (RBAC) availability: Checks that the invoking user's role is permitted to invoke the command. If the command modifies system state, the user's role must have modify permission for the command. If the command only displays system state, the user's role must have observe permission for the command. Some commands both observe and modify system state and thus require observe-modify permission. The following RBAC permissions are supported:
 - O = Observe
 - OM = Observe-modify,
 - N = None or not available
3. Virtual Fabric availability: If Virtual Fabrics are enabled, commands are checked for context and switch type as follows:
 - Virtual Fabric context (VF) = Command applies to the current logical switch only, or to a specified logical switch.
Virtual Fabric commands are further constrained by one of the following switch types:
 - All Switches (All) = Command can be run in any switch context.
 - Base Switch (BS) = Command can be run only on the base switch.
 - Default Switch (DS) = Command can be run only on the default switch.
 - N/A = Switch type is not applicable to the command.
 - Chassis context (CH) = Command applies to the chassis on which it is executed.
 - Switch and Chassis context (VF/CH) = Command applies to the switch and the chassis.
 - Disallowed = Command cannot be executed when Virtual Fabrics are enabled.
4. Command-specific: Checks whether the command is supported on the platform for which it is targeted.
5. Remote Domain Execution: Commands can be executed from a specified remote domain or logical switch context. In most cases, the Show commands and the commands with show options are supported.

NOTE

To determine RBAC permissions for a specified command and associated major options, use the `classConfig --showcli` command. Refer to "Using Fabric OS Commands" for more information.

Command Virtual Fabrics and Remote Domain execution restrictions

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands.

| Command name | Context | Switch type | Remote Domain Execution |
|------------------|---------|-------------|-------------------------|
| aaaConfig | CH | N/A | Disallowed |
| ad | VF | N/A | Disallowed |
| ag | VF | All | Disallowed |
| agAutoMapBalance | VF | All | Disallowed |
| agShow | VF | All | Disallowed |
| aliAdd | VF | All | Disallowed |
| aliCreate | VF | All | Disallowed |
| aliDelete | VF | All | Disallowed |
| aliRemove | VF | All | Disallowed |
| aliShow | VF | All | Disallowed |
| appLoginHistory | VF | All | Disallowed |
| aptPolicy | VF | All | Disallowed |
| auditCfg | CH | N/A | Disallowed |
| auditDump | CH | N/A | Disallowed |
| authUtil | VF | All | Disallowed |
| bannerSet | VF | N/A | Disallowed |
| bannerShow | VF | N/A | Disallowed |
| bcastShow | VF | All | Allowed |
| bladeCfgGemode | CH | N/A | Allowed |
| bladeDisable | CH | N/A | Disallowed |
| bladeEnable | CH | N/A | Disallowed |
| bladeSwap | VF | All | Disallowed |
| bootLunCfg | VF | All | Disallowed |
| bufOpMode | CH | N/A | Allowed |
| cfgActvShow | VF | All | Disallowed |
| cfgAdd | VF | All | Disallowed |
| cfgClear | VF | All | Disallowed |
| cfgCreate | VF | All | Disallowed |
| cfgDelete | VF | All | Disallowed |
| cfgDisable | VF | All | Disallowed |
| cfgEnable | VF | All | Disallowed |
| cfgRemove | VF | All | Disallowed |
| cfgSave | VF | All | Disallowed |
| cfgShow | VF | All | Disallowed |
| cfgSize | VF | All | Disallowed |
| cfgTransAbort | VF | All | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|-------------------------|------------|-------------|-------------------------|
| cfgTransShow | VF | All | Disallowed |
| chassisBeacon | CH | N/A | Disallowed |
| chassisCfgPerrThreshold | CH | N/A | Allowed |
| chassisDisable | CH | N/A | Disallowed |
| chassisDistribute | CH | N/A | Disallowed |
| chassisEnable | CH | N/A | Disallowed |
| chassisName | CH | N/A | Disallowed |
| chassisShow | CH | N/A | Allowed |
| classConfig | CH | N/A | Disallowed |
| cliHistory | CH/ VF | All | Disallowed |
| configDefault | VF/CH | N/A | Disallowed |
| configDownload | VF/CH | N/A | Disallowed |
| configList | CH | N/A | Disallowed |
| configRemove | VF | N/A | Disallowed |
| configShow | VF/CH | N/A | Disallowed |
| configUpload | VF/CH | N/A | Disallowed |
| configure | VF | All | Disallowed |
| configureChassis | CH | N/A | Disallowed |
| creditRecovMode | CH | N/A | Allowed |
| dataTypeShow | CH | N/A | Disallowed |
| date | CH | N/A | Disallowed |
| dbgShow | CH | N/A | Disallowed |
| defZone | VF | All | Disallowed |
| diagClearError | CH | N/A | Disallowed |
| diagDisablePost | CH | N/A | Disallowed |
| diagEnablePost | CH | N/A | Disallowed |
| diagHelp | Disallowed | N/A | Disallowed |
| diagPost | CH | N/A | Disallowed |
| diagRetry | CH | N/A | Disallowed |
| diagShow | VF | All | Disallowed |
| distribute | VF | All | Disallowed |
| dlsReset | VF | All | Disallowed |
| dlsSet | VF | All | Disallowed |
| dlsShow | VF | All | Allowed |
| dnsConfig | CH | N/A | Disallowed |
| enclosureShow | CH | N/A | Disallowed |
| errClear | CH | N/A | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|---------------------|------------|-------------|-------------------------|
| errDelimiterSet | CH | N/A | Disallowed |
| errDump | VF/CH | N/A | Disallowed |
| errFilterSet | CH | N/A | Disallowed |
| errModuleShow | VF | N/A | Disallowed |
| errShow | VF/CH | N/A | Disallowed |
| ethIf | CH | N/A | Allowed |
| exit | Disallowed | N/A | Disallowed |
| extnCfgr | VF | All | Disallowed |
| fabRetryShow | VF | All | Allowed |
| fabRetryStats | VF | All | Allowed |
| fabricLog | VF | All | Disallowed |
| fabricName | VF | All | Allowed |
| fabricPrincipal | VF | All | Allowed |
| fabricShow | VF | All | Allowed |
| fabStatsShow | VF | All | Allowed |
| fanDisable | CH | N/A | Disallowed |
| fanEnable | CH | N/A | Disallowed |
| fanShow | CH | N/A | Allowed |
| faPwn | VF | All | Allowed |
| fastBoot | CH | N/A | Disallowed |
| fcipHelp | Disallowed | N/A | Disallowed |
| fcipLedTest | CH | N/A | Disallowed |
| fcipPathTest | CH | N/A | Disallowed |
| fcPing | VF | All | Disallowed |
| fcplLogClear | VF | All | Disallowed |
| fcplLogDisable | VF | All | Disallowed |
| fcplLogEnable | VF | All | Disallowed |
| fcplLogShow | VF | All | Disallowed |
| fcplProbeShow | VF | All | Allowed |
| fcplRIsProbe | VF | All | Allowed |
| fcplRIsShow | VF | All | Allowed |
| fcrBcastConfig | VF | BS | Allowed |
| fcrConfigure | VF | All | Disallowed |
| fcrEdgeShow | VF | BS | Allowed |
| fcrFabricShow | VF | BS | Allowed |
| fcrIclPathBWMonitor | VF | All | Disallowed |
| fcrLsan | VF | BS | Allowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|------------------------|------------|-------------|-------------------------|
| fcrLsanCount | VF | BS | Allowed |
| fcrLsanMatrix | VF | BS | Allowed |
| fcrPhydevShow | VF | BS | Allowed |
| fcrProxyConfig | VF | BS | Disallowed |
| fcrProxyDevShow | VF | BS | Allowed |
| fcrResourceShow | VF | BS | Allowed |
| fcrRouterPortCost | VF | BS | Allowed |
| fcrRouteShow | VF | BS | Allowed |
| fcrXlateConfig | VF | BS | Allowed |
| fddCfg | VF | All | Disallowed |
| fdmiCacheShow | VF | All | Allowed |
| fdmiShow | VF | All | Allowed |
| femDump | CH | N/A | Allowed |
| ficonCfg | VF | All | Disallowed |
| ficonClear | VF | All | Disallowed |
| ficonCupSet | VF | All | Disallowed |
| ficonCupShow | VF | All | Allowed |
| ficonHelp | Disallowed | N/A | Disallowed |
| ficonShow | VF | All | Allowed |
| fipsCfg | CH | NA/ALL | Disallowed |
| firmwareActivate | CH | N/A | Disallowed |
| firmwareCheck | CH | N/A | Disallowed |
| firmwareCleanInstall | CH | N/A | Disallowed |
| firmwareCommit | CH | N/A | Disallowed |
| firmwareDownload | CH | N/A | Disallowed |
| firmwareDownloadStatus | CH | N/A | Allowed |
| firmwareKeyShow | CH | N/A | Allowed |
| firmwareRestore | CH | N/A | Disallowed |
| firmwareShow | VF | N/A | Allowed |
| firmwareSync | CH | N/A | Disallowed |
| flow | VF | All | Allowed |
| fosConfig | CH | N/A | Disallowed |
| fosExec | VF | All | Disallowed |
| frameLog | VF | All | Allowed |
| fspfShow | VF | All | Allowed |
| gePortErrShow | VF | All | Allowed |
| h | Disallowed | N/A | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|-----------------|------------|-------------|-------------------------|
| haDisable | CH | N/A | Disallowed |
| haDump | CH | N/A | Allowed |
| haEnable | CH | N/A | Disallowed |
| haFailover | CH | N/A | Disallowed |
| haReboot | CH | N/A | Disallowed |
| haRedundancy | CH | N/A | Allowed |
| haShow | CH | N/A | Allowed |
| haSyncStart | CH | N/A | Disallowed |
| haSyncStop | CH | N/A | Disallowed |
| help | Disallowed | N/A | Disallowed |
| historyLastShow | CH | N/A | Allowed |
| historyMode | CH | N/A | Disallowed |
| historyShow | CH | N/A | Allowed |
| i | CH | N/A | Disallowed |
| iflShow | VF | All | Allowed |
| interfaceShow | VF | All | Allowed |
| iodReset | VF | All | Disallowed |
| iodSet | VF | All | Allowed |
| iodShow | VF | All | Allowed |
| ipAddrSet | CH | N/A | Disallowed |
| ipAddrShow | VF, CH/ CH | N/A | Allowed |
| ipFilter | CH | N/A | Disallowed |
| ipSecConfig | CH | N/A | Disallowed |
| islShow | VF | All | Allowed |
| itemList | VF | All | Disallowed |
| killTelnet | CH | N/A | Disallowed |
| ldapCfg | CH | N/A | Disallowed |
| lfCfg | CH | N/A | Allowed |
| licenseAdd | CH | N/A | Disallowed |
| licenseIdShow | CH | N/A | Disallowed |
| licensePort | CH | N/A | Allowed |
| licenseRemove | CH | N/A | Disallowed |
| licenseShow | CH | N/A | Disallowed |
| licenseSlotCfg | CH | N/A | Disallowed |
| linkCost | VF | All | Disallowed |
| logicalGroup | VF | All | Allowed |
| login | Disallowed | N/A | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|--------------------|------------|-------------|-------------------------|
| logOut | Disallowed | N/A | Disallowed |
| lsanZoneShow | VF | BS | Disallowed |
| lsCfg | CH | N/A | Allowed |
| lsDbShow | VF | All | Allowed |
| mapsConfig | VF | All | Allowed |
| mapsDb | VF | All | Allowed |
| mapshelp | VF | All | Disallowed |
| mapsPolicy | VF | All | Allowed |
| mapsRule | VF | All | Allowed |
| mapsSam | VF | All | Allowed |
| memShow | CH | N/A | Disallowed |
| motd | CH | N/A | Disallowed |
| msCapabilityShow | VF | All | Disallowed |
| msConfigure | VF | All | Disallowed |
| msPlatShow | VF | All | Disallowed |
| msPlatShowDBCB | VF | All | Disallowed |
| msPIClearDB | VF | All | Disallowed |
| msPIMgmtActivate | VF | All | Disallowed |
| msPIMgmtDeactivate | VF | All | Disallowed |
| msTdDisable | VF | All | Disallowed |
| msTdEnable | VF | All | Disallowed |
| msTdReadConfig | VF | All | Disallowed |
| myld | VF | N/A | Disallowed |
| nbrStateShow | VF | All | Allowed |
| nbrStatsClear | VF | All | Disallowed |
| nodeFind | VF | All | Disallowed |
| nsAliasShow | VF | All | Disallowed |
| nsAllShow | VF | All | Disallowed |
| nsCamShow | VF | All | Disallowed |
| nsDevLog | VF | All | Disallowed |
| nsShow | VF | All | Disallowed |
| nsZoneMember | VF | All | Disallowed |
| nsZoneShow | VF | All | Disallowed |
| passwd | VF/CH | N/A | Disallowed |
| passwdCfg | CH | N/A | Disallowed |
| pathInfo | VF | All | Disallowed |
| pdShow | CH | N/A | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|--------------------------|---------|-------------|-------------------------|
| portAddress | VF | All | Allowed |
| portAlpaShow | VF | All | Allowed |
| portBeacon | VF | All | Disallowed |
| portBufferCalc | VF | All | Disallowed |
| portBufferShow | VF | All | Allowed |
| portCamShow | VF | All | Allowed |
| portCfg | VF | All | Disallowed |
| portCfgAlpa | VF | All | Disallowed |
| portCfgAutoDisable | VF | All | Allowed |
| portCfgCompress | VF | All | Allowed |
| portCfgCreditRecovery | VF | All | Allowed |
| portCfgDefault | VF | All | Disallowed |
| portCfgDPort | VF | All | Disallowed |
| portCfgEncrypt | VF | All | Disallowed |
| portCfgEport | VF | All | Disallowed |
| portCfgEportCredits | VF | All | Allowed |
| portcfgEXPort | VF | All | Disallowed |
| portCfgFaultDelay | VF | All | Disallowed |
| portCfgFec | VF | All | Allowed |
| portCfgfFillword | VF | All | Disallowed |
| portCfgFlogiLogout | VF | All | Disallowed |
| portCfgFportBuffers | VF | All | Disallowed |
| portCfgGE | VF | All | Allowed |
| portCfgGeMediaType | VF | All | Disallowed |
| portCfgGport | VF | All | Disallowed |
| portCfgISLMode | VF | All | Disallowed |
| portCfgLongDistance | VF | All | Disallowed |
| portCfgLossTov | VF | All | Disallowed |
| portCfgLport | VF | All | Disallowed |
| portCfgNonDfe | VF | All | Allowed |
| portCfgNPiVPort | VF | All | Disallowed |
| portCfgNPort | VF | All | Disallowed |
| portCfgOctetSpeedCombo | CH | N/A | Disallowed |
| portCfgPersistence | VF | All | Disallowed |
| portCfgPersistentDisable | VF | All | Disallowed |
| portCfgPersistentEnable | VF | All | Disallowed |
| portCfgQos | VF | All | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|--------------------|---------|-------------|-------------------------|
| portCfgShow | VF | All | Allowed |
| portCfgSpeed | VF | All | Disallowed |
| portCfgTdz | VF | All | Disallowed |
| portCfgTrunkPort | VF | All | Disallowed |
| portCfgVEXPort | VF | All | Disallowed |
| portChannelShow | VF | All | Disallowed |
| portCmd | VF | All | Disallowed |
| portDebug | CH | N/A | Disallowed |
| portDecom | VF | All | Disallowed |
| portDisable | VF | All | Disallowed |
| portDPortTest | VF | All | Allowed |
| portEnable | VF | All | Disallowed |
| portEncCompShow | VF | All | Allowed |
| portErrShow | VF | All | Allowed |
| portFlagsShow | VF | All | Allowed |
| portLedTest | CH | N/A | Disallowed |
| portLogClear | VF | All | Disallowed |
| portLogConfigShow | VF | All | Disallowed |
| portLogDisable | VF | All | Disallowed |
| portLogDump | VF | All | Disallowed |
| portLogDumpPort | VF | All | Disallowed |
| portLogEnable | VF | All | Disallowed |
| portLogEventShow | VF | All | Disallowed |
| portLoginShow | VF | All | Disallowed |
| portLogPdisc | VF | All | Disallowed |
| portLogReset | VF | All | Disallowed |
| portLogResize | VF | All | Disallowed |
| portLogShow | VF | All | Disallowed |
| portLogShowPort | VF | All | Disallowed |
| portLogTypeDisable | VF | All | Disallowed |
| portLogTypeEnable | VF | All | Disallowed |
| portLoopbackTest | CH | N/A | Disallowed |
| portName | VF | All | Disallowed |
| portPerfShow | VF | All | Disallowed |
| portPeerBeacon | VF | All | Allowed |
| portRouteShow | VF | All | Allowed |
| portShow | VF | All | Allowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|-------------------|------------|-------------|-------------------------|
| portStats64Show | VF | All | Allowed |
| portStatsClear | VF | All | Disallowed |
| portStatsShow | VF | All | Allowed |
| portSwap | VF | All | Disallowed |
| portSwapDisable | VF | All | Disallowed |
| portSwapEnable | VF | All | Disallowed |
| portSwapShow | VF | All | Allowed |
| portTest | VF | All | Disallowed |
| portTestShow | VF | All | Disallowed |
| portTrunkArea | VF | All | Allowed |
| portZoneShow | VF | All | Disallowed |
| powerOffListSet | CH | N/A | Disallowed |
| powerOffListShow | CH | N/A | Allowed |
| psShow | CH | N/A | Allowed |
| rasAdmin | CH | N/A | Disallowed |
| rasMan | Disallowed | N/A | Disallowed |
| reBoot | CH | N/A | Disallowed |
| relayConfig | CH | N/A | Allowed |
| roleConfig | CH | N/A | Disallowed |
| routeHelp | Disallowed | N/A | Disallowed |
| rtLogTrace | CH | N/A | Disallowed |
| sddQuarantine | VF | All | Disallowed |
| secActiveSize | VF | All | Disallowed |
| secAuthSecret | VF | All | Disallowed |
| secCertUtil | CH | N/A | Disallowed |
| secCryptoCfg | CH | N/A | Disallowed |
| secDefineSize | VF | All | Disallowed |
| secGlobalShow | VF | All | Disallowed |
| secHelp | Disallowed | N/A | Disallowed |
| secPolicyAbort | VF | All | Disallowed |
| secPolicyActivate | VF | All | Disallowed |
| secPolicyAdd | VF | All | Disallowed |
| secPolicyCreate | VF | All | Disallowed |
| secPolicyDelete | VF | All | Disallowed |
| secPolicyDump | VF | All | Disallowed |
| secPolicyFCSMove | VF | All | Disallowed |
| secPolicyRemove | VF | All | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|----------------------------|---------|-------------|-------------------------|
| secPolicySave | VF | All | Disallowed |
| secPolicyShow | VF | All | Disallowed |
| secStatsReset | VF | All | Disallowed |
| secStatsShow | VF | All | Disallowed |
| sensorShow | CH | N/A | Allowed |
| serDesTuneMode | CH | N/A | Allowed |
| setContext | VF | All | Disallowed |
| setDbg | CH | N/A | Disallowed |
| setVerbose | CH | N/A | Disallowed |
| sfpShow | VF | All | Disallowed |
| shellFlowControlDisable | CH | N/A | Disallowed |
| shellFlowControlEnable | CH | N/A | Disallowed |
| slotPowerOff | CH | N/A | Disallowed |
| slotPowerOn | CH | N/A | Disallowed |
| slotShow | VF, CH | N/A | Allowed |
| snmpConfig | CH | N/A | Disallowed |
| snmpTraps | CH | N/A | Disallowed |
| spinFab | VF | All | Disallowed |
| sshUtil | CH | N/A | Disallowed |
| statsClear | VF | All | Disallowed |
| stopPortTest | VF | DS | Disallowed |
| supportFfdc | CH | N/A | Disallowed |
| supportFtp | CH | N/A | Disallowed |
| supportInfoClear | CH | N/A | Allowed |
| supportSave | CH | N/A | Allowed |
| supportShow | VF | All | Disallowed |
| supportShowCfgDisable | CH | N/A | Disallowed |
| supportShowCfgEnable | CH | N/A | Disallowed |
| supportShowCfgShow | CH | N/A | Disallowed |
| switchBeacon | VF | All | Disallowed |
| switchCfgPersistentDisable | VF | All | Disallowed |
| switchCfgPersistentEnable | VF | All | Disallowed |
| switchCfgSpeed | VF | All | Disallowed |
| switchCfgTrunk | VF | All | Disallowed |
| switchDisable | VF | All | Disallowed |
| switchEnable | VF | All | Disallowed |
| switchName | VF | All | Disallowed |

TABLE 1 Virtual Fabric and Remote Domain execution restrictions for Fabric OS commands. (Continued)

| Command name | Context | Switch type | Remote Domain Execution |
|-------------------|------------|-------------|-------------------------|
| switchShow | VF | All | Allowed |
| switchUptime | CH | N/A | Allowed |
| switchViolation | VF | All | Disallowed |
| syslogAdmin | CH | N/A | Disallowed |
| sysShutdown | CH | N/A | Disallowed |
| tempShow | CH | N/A | Allowed |
| timeOut | CH | N/A | Disallowed |
| topologyShow | VF | All | Allowed |
| traceDump | CH | N/A | Disallowed |
| trunkDebug | VF | All | Allowed |
| trunkShow | VF | All | Allowed |
| tsClockServer | CH | N/A | Disallowed |
| tsTimeZone | CH | N/A | Disallowed |
| turboRamTest | CH | N/A | Disallowed |
| upTime | CH | N/A | Disallowed |
| uRouteShow | VF | All | Allowed |
| usbStorage | CH | N/A | Disallowed |
| userConfig | CH/VF | All/NA | Disallowed |
| version | VF | N/A | Allowed |
| wwn | VF, CH/ VF | All | Disallowed |
| wwnAddress | VF | All | Disallowed |
| wwnRecover | VF, CH/ VF | All | Disallowed |
| zone | VF | All | Disallowed |
| zoneAdd | VF | All | Disallowed |
| zoneCreate | VF | All | Disallowed |
| zoneDelete | VF | All | Disallowed |
| zoneHelp | Disallowed | N/A | Disallowed |
| zoneObjectCopy | VF | All | Disallowed |
| zoneObjectExpunge | VF | All | Disallowed |
| zoneObjectRename | VF | All | Disallowed |
| zoneObjectReplace | VF | All | Disallowed |
| zoneRemove | VF | All | Disallowed |
| zoneShow | VF | All | Disallowed |