
VRTX Chassis Diagnostics

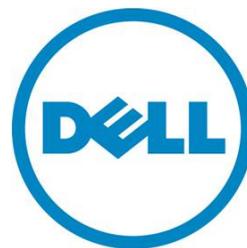
This white paper helps administrators to diagnose and troubleshoot the VRTX Chassis for any abnormalities in the modules listed (Power, PCIe Adapters, Servers, CMC, and Storage Components). VRTX Chassis logs the events on the different Logging Mechanisms such as SEL, LCD, and Chassis Log. It focusses on troubleshooting the events through the diagnostic commands and other troubleshooting techniques.

Author(s)

Anto Jesurajan

Arun Muthaiyan

Michael Brundridge



Executive summary

VRTX chassis has different logging mechanism to capture all the events on the chassis. The event could be configuration change or any critical or non-critical. The VRTX Chassis has different logging mechanism like Chassis log, SEL Log, LCD, Remote System Logging. The events could be identified from the logs and the recommended action can be performed as stated on the logs, especially on the Chassis logs to overcome the critical events.

Apart from the recommended action on the logs, Administrators can perform some diagnostic and recovery mechanism to overcome events. VRTX Chassis has a diagnostic console which could be used for identifying any network related issues or any other similar critical events. The different LED patterns on the CMC helps Administrators in troubleshooting the activity of the Chassis Controller. LEDs located on other components such as PSU, PCIe Adapter, and Storage Controller assists in their diagnoses.

Contents

Introduction	3
Terminology	3
Diagnostic Console	3
Traceroute	3
ping	4
ifconfig	4
gettracelog	4
racadm	5
Troubleshooting VRTX Components	7
Troubleshooting Power Supply Modules	7
Troubleshooting Fan Modules	8
Troubleshooting I/O Modules	9
Resetting Components	9
Restoring VRTX Chassis Configuration	10
Troubleshooting Network Time Protocol (NTP) Errors	10
Interpreting LED Colors and Blinking Patterns	11
Troubleshooting Non-responsive CMC	12
Observing LEDs to Isolate the Problem.....	12
Recovering Firmware Image	12
Racdump	13
Power Troubleshooting	13
Troubleshooting PCIe Components	15
Troubleshooting Storage Components	16
Conclusion	17



Introduction

This white paper focusses on troubleshooting various chassis events by Diagnostic Console, LED Pattern, Recover commands, Component troubleshooting such as PSU, PCIe Adapter, Storage Controller, and Chassis Controller. This could be helpful for administrators to try few troubleshooting techniques during critical events.

Terminology

RACADM: Administrator tool for configuring VRTX Chassis and servers

SEL: System event log a Logging standard from IPMI Specification

LCD: A liquid-crystal display is a flat panel display

NTP: Network Time Protocol

PSU: Power Supply Unit

LAN: Local Area Network

CMC: Chassis Management Controller

IPV6: Internet Protocol Version 6

Diagnostic Console

The VRTX Chassis troubleshooting can be done through the chassis Diagnostic commands. Issues related to the chassis hardware can be diagnosed using the RACADM commands, if logged in as an advanced user. To modify these settings, you must have the **Debug Command Administrator** privilege. To access the Diagnostic Console using the CMC Web interface, do the following:

In the VRTX CMC Web interface, click Chassis Overview > **Troubleshooting** > **Diagnostics**. In the **Command** text box, type a command and click **Submit** to obtain the response. The Diagnostic commands explained here can be run on the console and the corresponding response can be obtained.

Traceroute

Traceroute is a Linux-standard command to trace the path of the packets, traverse time in the network across Internet protocol. The following is the sample command and response:

```
$ traceroute 10.36.0.26
traceroute to 10.36.0.26 (10.36.0.26), 30 hops max, 40 byte packets
10.36.0.26 (10.36.0.26) 0.221 ms 0.149 ms 0.25 ms
```

The mean time in each hop can be added to obtain the overall time it took to reach the host. If the response comes with a "*" in any of the fields, then it implies that the data cannot be fetched and it could be various reason such as lookup failure, any network packet loss, or unclear routing. If the response has this type of situation on the response, then the DNS and routing on the network has to be diagnosed.

This command can be used on the diagnostic console to detect any issues in the network path on the VRTX Chassis. Any configurations made on the VRTX chassis for NTP servers, Remote syslog IPs, and other remote management stations in the network can be diagnosed. Traceroute6 is the equivalent command used on the IPV6 network.

For example:

```
traceroute6 <IP Address or Host name>
```



ping

The `ping` command is used to check the connectivity between systems in the network. When a Web connectivity is unable to connect to the VRTX Server on the Internet or LAN, the cause is often, because the Web server is not functioning, or it may be a network-related issue which hinders from connecting the host system. Therefore, the first step in diagnosing the problem is to test if the network connection is working. The `ping` command can be used in this situation without requiring a Web server. In the diagnostic console "ping" followed by the URL or IP address of the host you want to test. `Ping6` is the equivalent command to IPV6 networks. For example, `ping < IP Address or Hostname >`

```
The following is the sample command and response: $ ping 127.0.0.1
PING 127.0.0.1 (127.0.0.1): 56 data bytes
64 bytes from 127.0.0.1: icmp_seq=0 ttl=64 time=0.1 ms
--- 127.0.0.1 ping statistics ---
1 packets transmitted, 1 packets received, 0 percent packet loss
```

From the above command, the delay time and the packet loss can be diagnosed.

ifconfig

`ifconfig` -(Interface configurator) is widely used to initialize the network interface and to enable or disable the interfaces. In the VRTX Chassis `ifconfig` command, when issued on the **Diagnostic** console, displays the configuration of the active interface on the chassis to the external network. This helps to diagnose any issues on the VRTX chassis interface settings.

For Instance: `ifconfig`

```
eth0      Link encap:Ethernet  HWaddr 78:45:C4:FA:31:90
          inet addr:10.36.0.26  Bcast:10.36.1.255  Mask:255.255.254.0
          inet6 addr: fe80::7a45:c4ff:fefa:3190/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:13180720 errors:0 dropped:0 overruns:0 frame:0
          TX packets:489876 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1687604342 (1609.4 Mb)  TX bytes:54119040 (51.6 Mb)
```

Where `eth0` is the interface

`inet addr` is the IP address.

`Mask` is the subnet mask.

`HWaddr` is the MAC Address of the device.

gettracelog

The `gettracelog` command gives access to logs which can help troubleshoot the issues on any settings or abnormal activities on the chassis, apart from the chassis log and SEL log. The `gettracelog` response helps identify the failure more in detail than the Chassis log and SEL log.



racadm

The racadm command is the Administration & Configuration command line utility on a VRTX Chassis. VRTX Chassis supports the following commands. Any help about a specific command can be obtained by running the command.

```
racadm help <sub command> .
```

The syntax of the command is: racadm <command> <args>

List of supported Racadm Commands on VRTX Chassis is given here.

```
help                -- list racadm subcommand description
help <subcommand>  -- display usage summary for a subcommand
?                  -- list racadm subcommand description
? <subcommand>    -- display usage summary for a subcommand
arp                -- display the networking arp table
chassisaction      -- execute power-up/down/cycle or reset operation
chassislog         -- display the chassislog
closessn          -- close a session
clrraclog         -- clear the CMC log
clrsel            -- clear the System Event Log (SEL)
cmcchangeover      -- changes the redundant state of the CMC from active to
                    standby and vice versa
config            -- modify CMC configuration properties
connect           -- connect to switch or blade serial console
deploy            -- deploy blade or IOM by specifying required properties
eventfilters      -- configure alerts for chassis events
fanoffset         -- override the fan speed for fan 1 - 6
feature           -- display features active on the chassis / feature
                    deactivation
featurecard        -- feature card status and list the available features
fwupdate          -- update the firmware on a CMC, server, IOM inf, PERC, or
                    HDD
get               -- save CMC Event Filter configuration to a file
getactiveerrors   -- display CMC active errors
getassettag       -- display asset tag
getchassisname    -- get the chassisname
getconfig         -- display CMC configuration properties
getdcinfo         -- display general I/O module and DC configuration
                    information
getflexaddr       -- display Flexaddress enablement status for all slots and
                    fabrics
getioinfo         -- display general IO information and stack information
getled            -- display the LED settings on a module
getmacaddress     -- get MAC/WWN addresses
getmodinfo        -- get module configuration and status information
getniccfg         -- display network settings for modules
getpbinfo         -- get power budget status information
getpciecfg        -- display pcie slot/adaptor information
getpminfo         -- get power management status information
getraclog         -- display the CMC log
getractime        -- display the current CMC time
getredundancymode -- gets the redundancy mode of the CMC
getsel           -- display records from the System Event Log (SEL)
```



```

getsensorinfo    -- display system sensors
getslotname      -- gets the name of the slot in the chassis
getssninfo       -- display session information
getsvctag        -- display service tag information
getsysinfo       -- display general CMC and system information
gettracelog      -- display the CMC diagnostic trace log
getversion       -- display version information for modules
ifconfig         -- display network interface information
jobqueue         -- Jobqueue of the jobs currently scheduled
krbkeytabupload  -- upload an Kerberos Keytab to the CMC
license          -- License Manager commands
netstat          -- display routing table and network statistics
ping             -- send ICMP echo packets on the network
ping6           -- send ICMP echo packets on the network
racdump          -- display CMC diagnostic information
racreset         -- perform a CMC or RAC reset operation
racresetcfg      -- restore the CMC configuration to factory defaults
racresetpcie     -- restore the CMC PCIE to factory defaults
raid             -- display and configure hardware RAID subsystem
remoteimage      -- connect, disconnect or deploy a media file on a remote
                  server
serveraction     -- perform system power management operations
set              -- import saved CMC Event Filter configuration from a file
setassettag      -- set the asset tag for the specified module
setchassisname   -- sets the name of the chassis
setflexaddr      -- enable/disable the Flexaddress feature on a per fabric,
                  per slot basis.
setled           -- set state of the LEDs on a module
setniccfg        -- modify network configuration properties
setpciecfg       -- configure pcie slot/adaptor assignment state
setractime       -- set the time on the CMC
setslotname      -- sets the name of the slot in the chassis
setsysinfo       -- set the chassis name and chassis location
sshpkauth        -- manage PK Authentication keys and accounts
sslcertdownload  -- download an SSL certificate from the CMC
sslcertupload    -- upload an SSL certificate to the CMC
sslcertview      -- display a CA/server certificate in the CMC
sslcsrgen        -- generate a certificate CSR from the CMC
sslresetcfg      -- generate a new self-signed certificate
testemail        -- test CMC e-mail notifications
testfeature      -- test CMC feature x
testtrap         -- test CMC SNMP trap notifications
traceroute       -- determine the route of a packet
traceroute6      -- determine the route of a packet
wsman            -- perform wsman client functions for servers

```

The `racadm help <command>` displays the syntax and the usage of the specific command given as an argument to the help command.



Troubleshooting VRTX Components

The following procedures describe how to troubleshoot the following components:

- Power supply modules
- Fan module
- CMC module
- Network switch module

Troubleshooting Power Supply Modules

The power supply modules are hot-pluggable. It is highly recommended to hot-plug one PSU at a time, because removal of two or more PSUs may cause the Chassis or the Server to automatically turn off on the basis of Power Supply Redundancy Configuration. The PSUs may take few minutes to initialize. Therefore, it is advised to hot-plug the PSU only one at a time. The CMC1.0 for PowerEdge VRTX supports 1100W PSUs, which require either 110V or 220V input level.

The PSUs have an AC indicator which is green when connected to the PSU. If the LED is not illuminated, it implies that AC supply is not available to the PSU. This may be a fault in the PSU, the power grid supplying the input power, or in the input cable connected to the PSUs.

After the PSU are installed, all the LEDs in the PSUs indicate a healthy PSU, and the chassis health shows no critical events, then the Chassis is ready for operation. If the chassis has issues turning on any modules or the chassis itself, check the Chassis Log, SEL Log, Chassis Log, or the LCD for further information to assist in troubleshooting.

The `racadm getmodinfo` command gets the health of the components in the VRTX Chassis. The following is an example of the command and sample output:

```
$ getmodinfo
<module>          <presence>      <pwrState>      <health>        <svcTag>
Chassis           Present         ON              OK              PLST005
Main-Board       Present         ON              OK              N/A
Storage          Present         ON              OK              PLST005
Fan-1            Present         ON              OK              N/A
Fan-2            Present         ON              OK              N/A
Fan-3            Present         ON              OK              N/A
Fan-4            Present         ON              OK              N/A
Fan-5            Present         ON              OK              N/A
Fan-6            Present         ON              OK              N/A
Blower-1         Present         ON              OK              N/A
Blower-2         Present         ON              OK              N/A
Blower-3         Present         ON              OK              N/A
Blower-4         Present         ON              OK              N/A
PS-1             Present         Online          OK              N/A
PS-2             Present         Online          OK              N/A
PS-3             Present         Online          OK              N/A
PS-4             Present         Online          OK              N/A
CMC-1            Present         Standby        OK              N/A
CMC-2            Present         Primary        OK              N/A
```



Switch-1	Present	ON	OK	N/A
Server-1	Present	ON	OK	CFGBLD4
Server-2	Present	OFF	OK	G1BPNW1
Server-3	Present	OFF	OK	N/A
Server-4	Present	OFF	OK	N/A
DVD	Not Present	N/A	N/A	N/A
IO-Cable	Present	ON	OK	PLST005
FPC-Cable	Present	ON	OK	PLST005

In the `getmodinfo` command response, the PS-1..PS-4 Module indicates the power state and health condition of the PSU. If any of the PSUs is not functioning, it may indicate as **Failed** in the **Power State**, **Critical** or **Not OK** in **Health**. In that case, service the PSU.

Troubleshooting Fan Modules

The Fan modules on VRTX Chassis are hot-pluggable devices. VRTX Chassis has two types of fans —Blowers and internal storage fans. It is advised to remove and replace only one fan module at a time in a system that is turned on. Removal of all the internal fans at one time may cause the chassis to overheat, resulting in a thermal fault and the turning off of various components. Each Fan module has indicators that identify a non-functioning fan, or by running the `getmodinfo RACADM` command. For any fan issues related to the fans, VRTX Chassis will log a message to the SEL Log or the Chassis Log.

From the sample output defined in the `Getmodinfo racadm` command response, fan 1—fan 6 and blower 1—blower 4 are the fans, the `powerstate` and `health` field reflect the state and health of a fan. If the health state of a fan is not okay, it is advised to remove the fan and replace it. If fan-3, which is the CMC fan, is not functioning, it may result in the generation of a Pfault in the SEL Log and the chassis not powering on. If the problem is not resolved after a chassis powercycle, replace the fan.

Internal Chassis Fans can be set to different fan speeds running the `racadm fanoffset` command in cases of the internal fan zone requiring more air flow.

- Disable the fan-offset feature
`racadm fanoffset -s off`
- Increase fan speed by 20% of fan's maximum speed.
`racadm fanoffset -s low`
- Increases fan speed by 50% of a fan's maximum speed.
`racadm fanoffset -s medium`
- Set fan to run at 100% of a fan's maximum speed
`racadm fanoffset -s high`



Troubleshooting I/O Modules

To eliminate the possibility of a hardware issue with the module or its attaching devices, make sure that the module is properly initialized and configured.

Make sure that you have installed the module in an I/O slot that matches its fabric type. Check that the pass-through or switch module is cabled correctly.

For switch modules (pass-through modules are non-managed), run the `Connect Switch` command to verify that the switch is fully started, and verify the firmware revision and IP address of the switch. Verify that the switch module has a valid IP address for the subnet. Verify using the `ICMP ping` command. Check the network connector indicators on the network switch module.

If the link indicator displays an error status, check all cable connections. Try another connector on the external switch or hub. If the activity indicator does not illuminate, replace the network switch module.

Using the switch management interface, verify the switch port properties. If the switch is configured correctly, back up the switch configuration and replace the switch. See the switch module documentation for details. If the network link indicator for a given server is green, then the server has a valid link to the appropriate network switch module.

Make sure that the appropriate operating system drivers are installed and that the protocol settings are configured to insure proper communications.

Resetting Components

VRTX Chassis CMC module can be reset using reset components, reset iDRAC without restarting the operating system, or to virtually reseal servers causing them to behave as if they were removed and reinserted. If the chassis has a standby CMC, resetting the active CMC causes a failover and the standby CMC becomes active.

To reset components, you must have the administrator privilege. To reset the components using the CMC Web interface, select **Chassis Overview > Troubleshooting > Reset Components**.

To reset the active CMC, in the **CMC Status** section, click **Reset/Failover CMC**. If a standby CMC is present and a chassis is fully redundant, a failover occurs causing the standby CMC to become active. The same can be achieved by running the RACADM command `racadm cmchangeover`.

To reset an iDRAC without restarting the server operating system, in the **Reset Server** section, click **iDRAC Reset** in the **Reset** drop-down menu for the servers, whose iDRAC you want to reset, and then click **Apply Selections**. This resets the iDRACs for the servers without restarting the operating system. To reset only an iDRAC without restarting the operating system using RACADM, see the *Chassis Management Controller for PowerEdge VRTX -RACADM Command Line Reference Guide*.

When an iDRAC is reset, fans are set to 100% for the server that may cause the blowers to increase in speed for a short period of time, which is normal. It is recommended to try resetting the iDRAC before you attempt to virtually reseal the servers.



To virtually reseal the server, in the **Reset Server** section, click **Virtual Reseat** in the **Reset** drop-down menu for the servers you want to reseal, and then click **Apply Selections**. This operation causes the servers to behave as if they were removed and reinserted.

The following RACADM command can also be used to reseal servers:

```
racadm serveraction -m <server-n> reseal -f  
where n is the server number.
```

Restoring VRTX Chassis Configuration

VRTX Chassis configuration can be backed up and restored in case of any issues in the chassis. To save or restore a backup of the Chassis configuration using the CMC Web interface, select **Chassis Overview > Setup > Chassis Backup**.

To save the chassis configuration, in the **Save** window, click **Save**. Override the default file path (optional) and click **OK** to save the file.

The default backup file name contains the service tag of the chassis. This backup file can be used later to restore the settings and certificates for this chassis only as the file is encrypted for use on the chassis generating the backup. To restore the chassis configuration, in the **Restore** window click **Browse**, specify the backup file, and then click **Restore**.

VRTX Chassis CMC does not reset after restoring the configuration. However, CMC services may take some time to effectively impose any changed or new configuration. After successful completion, all current sessions are closed.

Troubleshooting Network Time Protocol (NTP) Errors

After configuring CMC to synchronize its clock with a remote time server over the network, it may take 2–3 minutes before a change occurs in the date and time. After this time, if there is still no change, it may be necessary to troubleshoot an issue. CMC may not be able to synchronize its clock for the following reasons:

- Issue with the settings of NTP Server1, NTP Server2, and NTP Server3.
- Invalid host name or IP address may have been entered.
- Network connectivity issue that prevents CMC from communicating with any of the configured NTP servers.
- DNS problem, preventing any of the NTP server host names from being resolved.

To troubleshoot these issues, check the information in the CMC Trace Log. This log contains an error message for NTP-related issues. If CMC is not able to synchronize with any of the configured remote NTP servers, then CMC time is synchronized to the local system clock and the trace log contains an entry similar to the following:

```
Jan 8 20:02:40 cmc ntpd[1423]: synchronized to LOCAL(0), stratum 10
```

You can also check the ntpd status by running the following racadm command:

```
racadm gettractime -n
```



If the '*' is not displayed for one of the configured servers, the settings may not be configured correctly. The output of this command contains detailed NTP statistics that may be useful in debugging the issue. If attempting to configure a Windows-based NTP server, it may help to increase the MaxDist parameter for ntpd. Before changing this parameter, understand all the implications, because the default setting must be large enough to work with most NTP servers.

To modify the parameter, run the following command:

```
racadm config -g cfgRemoteHosts -o cfgRhostsNtpMaxDist 32
```

After changing, disable NTP, wait for 5–10 seconds, then enable NTP. NTP may take an additional three minutes to synchronize again.

To disable NTP, run the following command:

```
racadm config -g cfgRemoteHosts -o cfgRhostsNtpEnable 0
```

To enable NTP, run the following command:

```
racadm config -g cfgRemoteHosts -o cfgRhostsNtpEnable 1
```

If the NTP servers are configured correctly and this entry is present in the trace log, then this confirms that CMC is not able to synchronize with any of the configured NTP servers.

If the NTP server IP address is not configured, you may see a trace log entry similar to the following:

```
Jan 8 19:59:24 cmc ntpd[1423]: Cannot find existing interface for address 1.2.3.4
```

```
Jan 8 19:59:24 cmc ntpd[1423]: configuration of 1.2.3.4 failed
```

If an NTP server setting was configured with an invalid host name, you may see a trace log entry as follows:

```
Aug 21 14:34:27 cmc ntpd_initres[1298]: host name not found: error
```

```
Aug 21 14:34:27 cmc ntpd_initres[1298]: couldn't resolve hostname, giving up on it
```

Interpreting LED Colors and Blinking Patterns

LEDs on the chassis provide the component status as following:

- Steadily glowing green LEDs indicate that the component is turned on. If the green LED is blinking, it indicates a critical but routine event, such as a firmware upload, during which the unit is not operational. It does not indicate a fault.
- A blinking amber LED on a module indicates a fault on that module.
- Blue, blinking LEDs are configurable by the user and used for identification

Component	LED Blinking Pattern	Implication
CMC	Green Glowing steadily	Turned ON
	Green dark	Turned OFF
	Blue Glowing Steadily	Active
	Blue Dark	Standby
	Amber Blinking	Fault



Server	Green glowing Steadily	Turned on
	Green Dark	Turned off
	Blue, glowing steadily	Normal
	Blue, blinking	User-enabled module identifier
	Amber, glowing steadily	Not used
	Amber, blinking	Fault
	Blue, dark	No Fault
IOM (Common)	Green, glowing steadily	Powered on

Troubleshooting Non-responsive CMC

VRTX Chassis unresponsiveness of CMC using any of the interfaces (the Web interface, Telnet, SSH, remote RACADM, or serial), you can diagnose by observing the LEDs on CMC, obtaining recovery information using the DB-9 serial port, or recovering the CMC firmware image. It is not possible to log in to the standby CMC using a serial console.

Observing LEDs to Isolate the Problem

When the VRTX Chassis enclosure is opened, the Mainboard has two CMCs installed on the respective slots and there are two LEDs on the side of the card. Top green LED indicates power. If it is not turned on, verify that you have an AC power connected to at least one PSU.

Also make sure the CMC card is installed properly on the slots. Before the CMC is pulled out, make sure the Chassis is passive, and then remove CMC, reinstall CMC making sure the board is inserted all the way and that the latch closes correctly.

Bottom LED – the bottom LED is multi-colored. When CMC is active and running, and there are no problems, the bottom LED is blue. If it is amber, a fault was detected. The fault may be caused by any of the following:

- A CMC Core failure, the CMC board must be replaced.
- A CMC Self-test failure, the CMC board must be replaced.
- Corrupted CMC Firmware and Recovery image. An image corruption can be rectified by uploading the CMC firmware image to recover the CMC (see Recovering Firmware Image).

A normal CMC start or reset takes about a minute to fully boot into its operating system and be available for log in. The blue LED is illuminated on the active CMC. In a redundant, two-CMC configuration, only the top green LED is illuminated on the standby CMC.

Recovering Firmware Image

CMC enters recover mode when a normal CMC operating boot is not possible. To utilize recovery mode, connect a null modem cable to the VRTX Chassis serial port. In recover mode, a small subset of commands are available that allow you to reprogram the flash devices by uploading the firmware



update file `vrtx_cmc.bin`. This is the same firmware image file used for normal firmware updates. The recovery process displays its current activity and boots to the CMC OS upon completion. When you type `recover` on the recovery prompt, the `recover` reason and available sub-commands display.

An example of a `recover` sequence is given here.

```
recover getniccfg
recover setniccfg 192.168.0.120 255.255.255.0 192.168.0.1
recover ping 192.168.0.100
recover fwupdate -g -a 192.168.0.100
```

In `recover` mode, you cannot ping CMC normally, because there is no active network stack. The `recover ping <TFTP server IP>` command allows you to ping to the TFTP server to verify the LAN connection. You must run the `recover reset` command after running `setniccfg` on some systems.

Racdump

The `racdump` subcommand provides a single command to get comprehensive chassis status, configuration state information, and the historic event logs. The `racdump` subcommand displays the VRTX Chassis information such as General system/RAC information, CMC information, Chassis information, Session information, Sensor information, Firmware build information, Chassis logs, and SEL Logs.

From the `racdump` command, the Chassis Information can be gathered and can be used for troubleshooting.

Power Troubleshooting

The following information helps you to troubleshoot power supply and power-related issues:

Issue: Configured the Power Redundancy Policy to AC Redundancy, and a **Power Supply Redundancy Lost** event was opened.

- A. **Resolution:** This configuration requires at least one PSU in grid A (the top two slots) and one PSU in grid B (the bottom two slots) to be present and functional in the modular enclosure. Additionally, the capacity of each grid must be enough to support the total power allocations for the chassis to maintain AC redundancy. (For full AC Redundancy operation, make sure that a full PSU configuration of four PSUs is available.)
- B. **Resolution:** Check if all PSUs are properly connected to the two AC grids; PSUs in grid A must be connected to one AC grid, those in grid B must be connected to the other AC grid, and both AC grids need to be functional. AC Redundancy is lost when one of the AC grids is not functioning.

Problem: The PSU state is displayed as **Failed (No AC)**, even when an AC cord is connected and the power distribution unit is producing good AC output.

- A. **Resolution:** Check and replace the AC cord. Check and confirm that the power distribution unit providing power to the PSU is working as expected. If the issue still persists, call your service provider for a PSU replacement.
- B. **Resolution:** Make sure that the PSU is connected to the same voltage as the other PSUs. If CMC detects that a PSU operating at a different voltage, the PSU is turned off and marked **Failed**.



Problem: Dynamic Power Supply Engagement (DPSE) is enabled, but none of the PSUs display in the Standby state.

- A. **Resolution:** There is insufficient surplus power. One or more PSUs are moved into the Standby state only when the surplus power available in the enclosure exceeds the capacity of at least one PSU.
- B. **Resolution:** DPSE cannot be fully supported with the PSUs present in the enclosure. To check if this is the case, use the Web interface to disable the Dynamic Power Supply Engagement feature, and then enable the feature. A message is displayed if DPSE cannot be fully supported.

Problem: Inserted a new server into the enclosure with sufficient PSUs, but the server does not turn on.

- A. **Resolution:** Check the system input power cap setting. The power cap might be configured too low to allow any additional servers to be turned on.
- B. **Resolution:** Check that the server model, iDRAC version and BIOS are supported on the chassis in question.
- C. **Resolution:** Check the maximum power conservation setting. If this is set, servers are allowed to turn on. For more details, see the power configuration settings.
- D. **Resolution:** Check the server slot power priority of the slot associated with the newly-inserted server, and make sure it is not lower than any other server slot power priority.
- E. **Resolution:** While not power related, check that the fabric of the I/O mezzanine cards are correct for the chassis configuration. If the I/O mezzanine card does not match the chassis configuration, the server will not be allowed to turn on. Remove or replace any I/O mezzanine card that does not match the chassis configuration and try again.

Problem: A subset of servers lost power after an AC Grid failure, even when the chassis was operating in the AC Redundancy configuration with four PSUs.

- **Resolution:** This can occur if the PSUs are improperly connected to the redundant AC grids at the time the AC grid failure occurs. The AC Redundancy policy requires that the top two PSUs be connected to one AC Grid, and bottom two PSUs to be connected to the other AC Grid. If two PSUs are improperly connected, such as PSU3 and PSU4 are connected to the incorrect AC grids, an AC grid failure results in the loss of power to the least-priority servers.

Problem: The least-priority servers lost power after a PSU failure.

- **Resolution:** This is expected behavior if the enclosure power policy was configured to **No Redundancy**. To avoid a PSU failure in future, causing servers to turn off, make sure that the chassis has at least three PSUs and is configured for the Power Supply Redundancy policy to prevent PSU failure from impacting server operation.

Problem: Overall server performance decreases when the ambient temperature increases in the data center.

- **Resolution:** This can occur if the **System Input Power Cap** is configured to a value that results in an increased power requirement by fans having to be made up by reduction in the power allocation to the servers. User can increase the **System Input Power Cap** to a higher value that allow for additional power allocation to the fans without an impact on server performance.

Problem: Server is unable to turn on and have unsupported Server on the Chassis.

- **Resolution:** This is because the Server inserted in the chassis is not supported on the chassis, or the firmware version of the Server is not compatible with the VRTX Chassis. It is



advised to remove the server and insert a supported server on the Chassis or upgrade the server to a version of firmware supporting the VRTX chassis, and then insert it back to the VRTX Chassis.

Problem: Server is unable to turn on with Chassis Intrusion on Chassis log.

- **Resolution:** This may happen when the VRTX Chassis cover is open and the latch is not closed. It is advised to close the enclosure and latch it properly before a server turn-on command is issued.

Troubleshooting PCIe Components

The PCIe assignments to Servers can be achieved through various interfaces such as Web interface, RACADM, WS-MAN, and so on. In the CMC Web interface, **PCIe Overview > Properties**, and **PCIe Overview > Setup** helps to view and configure the PCIe adapters. For example, the assignments and viewing the assignments can be achieved in RACADM by running the following commands

To Assign PCIe Slots

```
racadm setpciecfg assign -c <FQDD> -i <server slot>
```

To Unassign PCIe Slots

```
racadm setpciecfg unassign -c <FQDD>
```

Valid Options:

```
-c           : Option to select a particular PCIe adapter or  
              Virtual Adapter  
-i           : Server slot number  
FQDD        : FQDD Key of selected PCIe slot or Virtual Adapter  
-----
```

To view the assignments, run the following command.

```
Racadm getpciecfg -a
```

The sample output of the command is

```
racadm getpciecfg -a  
<PCIe Slot#> <Name>           <Power State> <Server Slot Name>  
<Server Slot#>  
1           Empty           N/A           Unmapped  
N/A  
2           Empty           N/A           Unmapped  
N/A  
3           Empty           N/A           Unmapped  
N/A  
4           Empty           N/A           Unmapped  
N/A  
5           Empty           N/A           Unmapped  
N/A  
6           Empty           N/A           Unmapped  
N/A
```



7	Empty	N/A	Unmapped
N/A			
8	Empty	N/A	Unmapped
N/A			

<PCIe Slot#>	<Virtual Adapter#>	<Server Slot Name>	<Server Slot#>
9	1	cfg8s1mojo.bc.lab	1
9	2	SLOT-02	2
9	3	SLOT-03	3
9	4	SLOT-04	4
10	1	cfg8s1mojo.bc.lab	1
10	2	SLOT-02	2
10	3	SLOT-03	3
10	4	SLOT-04	4

From the above command, administrators can view which PCIe slots are assigned to which server and the power state of the PCIe adapter.

The server that is turned on may lead to the turning-on of the PCIe slots assigned to the server. The corresponding PCIe slot power state can be obtained from the `getpciecfg` command with `-a` as a parameter. If the power state is not ON or the unassignment is failed, then the Chassis log should have the cause for the failure and the recommended action on the log can be performed to overcome the state.

A server must be turned off before assigning a PCIe adapter to that server.

Troubleshooting Storage Components

The storage subsystem has the following components

- Storage Controller (PERC)
- SAS Expander
- Storage Backplane
- HDDs

If the storage subsystem is not functioning, then check for the health status of the storage subsystem, which can be obtained by running the RACADM command `getmodinfo`, and then check whether or not the cable between the SAS Expander and the Mainboard is connected. Cabling errors are logged in the VRTX Chassis logs. The following RACADM commands will be useful for obtaining the Controller and the root node status.

To get the FQDD of RAID controllers
`racadm raid get controllers`

To get the status of the Root Node
`racadm raid get status`

To get the controller status



```
racadm raid get controllers:<FQDD> -p status
```

FQDD can be obtained from the command above. A sample FQDD is
RAID.ChassisIntegrated.1-1

Therefore the command is

```
racadm raid get controllers:RAID.ChassisIntegrated.1-1 -p status
```

To get the controller Logs

```
racadm raid exportlog:<controller FQDD> -l <CIFS or NFS share> -u <username>  
-p <password> [-f <filename>]
```

Conclusion

This white paper is useful for administrators who want to troubleshoot and diagnose events on the VRTX chassis, or wish to gather information before contacting their technical support representative.

Learn more

Visit Dell.com/PowerEdge for more information on Dell's enterprise-class servers.

© 2013 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell and the Dell logo are trademarks of Dell Inc. Microsoft, Windows, and the Windows logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

June 2013 | Rev 1.0

