Dell Chassis Management Controller Version 3.0 for Dell EMC PowerEdge VRTX

User's Guide



Notes, cautions, and warnings

i NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

MARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Copyright © 2017 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

Chapter 1: Overview	13
What is new in this release	14
Key Features	14
Management Features	14
Security Features	15
Chassis Overview	15
Minimum CMC Version	18
Supported Remote Access Connections	18
Supported Platforms	18
Supported Web Browsers	19
Managing Licenses	19
Types of Licenses	19
Acquiring Licenses	19
License Operations	19
License Component State or Condition and Available Operations	20
Managing Licenses Using CMC Web Interface	20
Managing Licenses Using RACADM	20
Licensable Features in CMC	21
Viewing Localized Versions of the CMC Web Interface	22
Supported Management Console Applications	22
How to use this Guide	22
Other Documents You May Need	23
Accessing support content from the Dell EMC support site	23
Chapter 2: Installing and Setting Up CMC	25
Before You Begin	
Installing CMC Hardware	
Checklist To Set up Chassis	
Basic CMC Network Connection	26
Installing Remote Access Software on a Management Station	26
Installing RACADM on a Linux Management Station	
Uninstalling RACADM From a Linux Management Station	27
Configuring a Web Browser	27
Proxy Server	27
Microsoft Phishing Filter	28
Certificate Revocation List Fetching	28
Downloading Files From CMC With Internet Explorer	28
Enabling Animations In Internet Explorer	28
Setting Up Initial Access to CMC	28
Configuring Initial CMC Network	
Interfaces and Protocols to Access CMC	
Launching CMC Using Other Systems Management Tools	
Downloading and Updating CMC Firmware	33
Setting Chassis Physical Location and Chassis Name	33

ing
ing
ing

Viewing Currently Installed Firmware Versions	48
Viewing Currently Installed Firmware Versions Using CMC Web Interface	48
Viewing Currently Installed Firmware Versions Using RACADM	48
Updating the CMC Firmware	48
Signed CMC Firmware Image	49
Updating CMC and Mainboard Firmware	49
Updating CMC Firmware Using Web Interface	
Updating CMC firmware using RACADM	50
Updating Chassis Infrastructure Firmware	
Updating Chassis Infrastructure Firmware Using CMC Web Interface	51
Updating Chassis Infrastructure Firmware Using RACADM	51
Updating Server iDRAC Firmware	
Updating Server iDRAC Firmware Using Web Interface	51
Updating Server Component Firmware	
Server Component Update Sequence	53
Enabling Lifecycle Controller	54
Choosing Server Component Firmware Update Type Using CMC Web Interface	
Filtering Components for Firmware Updates	54
Viewing Firmware Inventory	
Viewing Firmware Inventory Using CMC Web Interface	
Viewing Firmware Inventory Using RACADM	56
Saving Chassis Inventory Report Using CMC Web Interface	56
Configuring Network Share Using CMC Web Interface	
Lifecycle Controller Job Operations	
Reinstalling Server Component Firmware	58
Rolling Back Server Component Firmware	
Rolling Back Server Component Firmware Using the CMC Web Interface	
Upgrading Server Component Firmware	
Upgrading Server Component Firmware From File Using CMC Web Interface	59
Server Component Single Click Update Using Network Share	
Pre-requisites for Using Network Share Update Mode	60
Upgrading Server Component Firmware From Network Share Using CMC Web Interface	60
Supported Firmware Versions for Server Component Update	61
Deleting Scheduled Server Component Firmware Jobs	62
Deleting Scheduled Server Component Firmware Jobs Using the Web Interface	62
Updating Storage Component Using CMC Web Interface	
Recovering iDRAC Firmware Using CMC	62
hapter 5: Viewing Chassis Information and Monitoring Chassis and Component Health	64
Viewing Chassis and Component Summaries	64
Chassis Graphics	65
Selected Component Information	66
Viewing Server Model Name and Service Tag	68
Viewing Chassis Summary	
Viewing Chassis Controller Information and Status	69
Viewing Information and Health Status of All Servers	69
Viewing Health Status and Information for Individual Server	69
Viewing Information and Health Status of the IOM	69
Viewing Information and Health Status of Fans	69
Configuring Fans	70

Viewing Front Panel Properties	71
Viewing KVM Information and Health Status	71
Viewing LCD Information and Health	71
Viewing Information and Health Status of Temperature Sensors	
Viewing Storage Capacity and Status of the Storage Components	
Chapter 6: Configuring CMC	
Viewing and Modifying CMC Network LAN Settings	
Viewing and Modifying CMC Network LAN Settings Using CMC Web Interface	74
Viewing and Modifying CMC Network LAN Settings Using RACADM	
Enabling the CMC Network Interface	74
Enabling or Disabling DHCP for the CMC Network Interface Address	
Enabling or Disabling DHCP for DNS IP Addresses	
Setting Static DNS IP addresses	
Configuring IPv4 and IPv6 DNS Settings	
Configuring Auto Negotiation, Duplex Mode, and Network Speed for IPv4 and IPv6	76
Setting the Maximum Transmission Unit for IPv4 and IPv6	76
Configuring CMC Network and Login Security Settings	77
Configuring IP Range Attributes Using CMC Web Interface	77
Configuring IP Range Attributes Using RACADM	77
Configuring Virtual LAN Tag Properties for CMC	78
Configuring Virtual LAN Tag Properties for CMC Using RACADM	78
Configuring Virtual LAN Tag Properties for CMC Using Web Interface	79
Federal Information Processing Standards	79
Enabling FIPS Mode Using CMC Web Interface	79
Enabling FIPS Mode Using RACADM	80
Disabling FIPS Mode	80
Configuring Services	80
Configuring Services Using CMC Web Interface	80
Configuring Services Using RACADM	81
Configuring CMC Extended Storage Card	81
Setting Up Chassis Group	81
Adding Members To Chassis Group	82
Removing a Member from the Leader	82
Disbanding a Chassis Group	82
Disabling an Individual Member at the Member Chassis	83
Accessing the Web page of a Member Chassis or Server	83
Propagating Leader Chassis Properties to Member Chassis	83
Server Inventory for MCM group	84
Saving Server Inventory Report	84
Chassis Group Inventory and Firmware Version	85
Viewing Chassis Group Inventory	85
Viewing Selected Chassis Inventory Using Web Interface	85
Viewing Selected Server Component Firmware Versions Using Web Interface	86
Chassis Configuration Profiles	86
Saving Chassis Configuration	86
Restoring Chassis Configuration Profile	87
Viewing Stored Chassis Configuration Profiles	87
Applying Chassis Configuration Profiles	87
Exporting Chassis Configuration Profiles	87

Editing Chassis Configuration Profiles	88
Deleting Chassis Configuration Profiles	88
Configuring Multiple CMCs Using RACADM	88
Creating a CMC Configuration File	89
Parsing Rules	89
Modifying the CMC IP Address	
Configuring Multiple CMCs through RACADM Using Chassis Configuration Profiles	
Exporting Chassis Configuration profiles	
Importing Chassis Configuration profiles	
Parsing Rules	
Viewing and Ending CMC Sessions	
Viewing and Ending CMC Sessions Using Web Interface	93
Viewing and Ending CMC Sessions Using RACADM	
Chapter 7: Configuring Servers	94
Configuring Slot Names	94
Configuring iDRAC Network Settings	95
Configuring iDRAC QuickDeploy Network Settings	95
Assigning QuickDeploy IP Address to Servers	97
Modifying iDRAC Network Settings for Individual Server iDRAC	97
Modifying iDRAC Network Settings Using RACADM	
Configuring iDRAC Virtual LAN Tag Settings	98
Configuring iDRAC Virtual LAN Tag Settings Using RACADM	98
Configuring iDRAC Virtual LAN Tag Settings Using Web Interface	98
Setting First Boot Device	98
Setting First Boot Device For Multiple Servers Using CMC Web Interface	99
Setting First Boot Device For Individual Server Using CMC Web Interface	99
Setting First Boot Device Using RACADM	100
Configuring Server FlexAddress	
Configuring Remote File Share	100
Configuring Profile Settings Using Server Configuration Replication	100
Accessing Server Profiles Page	101
Adding or Saving Profile	101
Applying Profile	102
Importing Profile	102
Exporting Profile	102
Editing Profile	103
Deleting Profile	
Viewing Profile Settings	
Viewing Stored Profile Settings	
Viewing Profile Log	
Completion Status And Troubleshooting	
Quick Deploy of Profiles	
Assigning Server Profiles to Slots	
Boot Identity Profiles	
Saving Boot Identity Profiles	
Applying Boot Identity Profiles	
Clearing Boot Identity Profiles	
Viewing Stored Boot Identity Profiles	
Importing Boot Identity Profiles	
= = = = = = = = = = = = = =	

107
107
108
108
108
108
109
109
110
111
111
111
111
112
112
113
116
116
119
119
119
120
121
121
122
122
124
126
132
133
133
134
135
135
135
136
136 136
136 136
136 136
130 137
137 137
137 137
137 137
10/
137
137 137

Configuring CMC SSO Or Smart Card Login For Active Directory Users Using RACADM	138
Chapter 11: Configuring CMC to Use Command Line Consoles	139
CMC Command Line Console Features	
CMC Command Line Interface Commands	
Using Telnet Console With CMC	
Using SSH With CMC	
Supported SSH Cryptography Schemes	
Configure Public Key Authentication Over SSH	
Configuring Terminal Emulation Software	
Configuring Linux Minicom	
Connecting to Servers or Input Output Module Using Connect Command	
Configuring the Managed Server BIOS for Serial Console Redirection	
Configuring Windows for Serial Console Redirection	
Configuring Linux for Server Serial Console Redirection During Boot	
Configuring Linux for Server Serial Console Redirection After Boot	146
Chapter 12: Using FlexAddress and FlexAdress Plus	149
About FlexAddress	
About FlexAddress Plus	150
Viewing FlexAddress Activation Status	150
Configuring FlexAddress	
Configuring FlexAddress for Chassis-Level Fabric and Slots	
Viewing World Wide Name or Media Access Control Addresses	
Fabric Configuration	
Viewing WWN or MAC Address Information	
Viewing Basic WWN or MAC Address Information Using Web Interface	
Viewing Advanced WWN or MAC Address Information Using Web Interface	
Viewing WWN or MAC Address Information Using RACADM	
Command Messages	
FlexAddress DELL SOFTWARE LICENSE AGREEMENT	
Chapter 13: Managing Fabrics	
Fresh Power-up Scenario	
Monitoring IOM Health	
Configuring Network Settings for IOM	
Configuring Network Settings for IOM Using CMC Web Interface	
Configuring Network Settings for IOM Using RACADM	
Managing Power Control Operation for IOMs	
Enabling or Disabling LED Blinking for IOMs	160
Chapter 14: Managing and Monitoring Power	161
Redundancy Policies	
Grid Redundancy Policy	
Power Supply Redundancy Policy	
Dynamic Power Supply Engagement	
Default Redundancy Configuration	
Grid Redundancy	
Power Supply Redundancy	164

Power Budgeting For Hardware Modules	164
Server Slot Power Priority Settings	165
Assigning Priority Levels To Servers	165
Assigning Priority Levels To Servers Using CMC Web Interface	165
Assigning Priority Levels To Servers Using RACADM	165
Viewing Power Consumption Status	166
Viewing Power Consumption Status Using CMC Web Interface	166
Viewing Power Consumption Status Using RACADM	166
AC Power Recovery	166
Viewing Power Budget Status Using CMC Web Interface	166
Viewing Power Budget Status Using RACADM	166
Redundancy Status and Overall Power Health	167
Power Management After PSU Failure	167
Power Management After Removing PSU	167
New Server Engagement Policy	167
Power Supply and Redundancy Policy Changes in System Event Log	168
Configuring power budget and redundancy	168
Power Conservation and Power Budget	169
Maximum Power Conservation Mode	169
Server Power Reduction to Maintain Power Budget	169
110V PSUs AC Operation	169
Remote Logging	170
External Power Management	170
Configuring Power Budget and Redundancy Using CMC Web Interface	
Configuring Power Budget and Redundancy Using RACADM	171
Executing Power Control Operations	172
Executing Power Control Operations on the Chassis	172
Executing Power Control Operations on the Chassis Using Web Interface	172
Executing Power Control Operations on the Chassis Using RACADM	
Executing Power Control Operations on a Server	
Executing Power Control Operations for Multiple Servers Using CMC Web Interface	
Executing Power Control Operations on the IOM	
Executing Power Control Operations on IOM Using CMC Web Interface	
Executing Power Control Operations on the IOM Using RACADM	174
01 . 45 M 01 0.	475
Chapter 15: Managing Chassis Storage Viewing Status of the Storage Components	
Viewing Status of the Storage Components	
Viewing Fault-tolerant Troubleshooting Information of SPERC Using CMC Web Interface	
Assigning Virtual Adapters To Slots Using CMC Web Interface	
Fault-Tolerance in Storage Controllers	
Security Key Mismatch	
Resolving Security Key Mismatch Using CMC Web Interface	
Viewing Controller Properties Using CMC Web Interface	
Viewing Controller Properties Using RACADM	
Importing or Clearing Foreign Configuration	
Configuring Storage Controller Settings	
Configuring Storage Controller Settings Using CMC Web Interface	
Configuring Storage Controller Settings Using RACADM	
Shared PERC Controllers	180

Enabling or Disabling RAID Controller Using CMC Web Interface	181
Enabling or Disabling RAID Controller Using RACADM	182
Enabling or disabling fault tolerance of external RAID controller using RACADM	182
Viewing Physical Disk Properties Using the CMC Web Interface	183
Viewing Physical Disk Drives Properties Using RACADM	183
Identifying Physical Disks and Virtual Disks	
Assigning Global Hot Spares Using CMC Web Interface	183
Assigning Global Hot Spares Using RACADM	183
Recovering Physical Disks	184
Viewing Virtual Disk Properties Using CMC Web Interface	184
Viewing Virtual Disk Properties Using RACADM	184
Creating Virtual Disk Using CMC Web Interface	184
Managing Encryption Keys	185
Creating Encryption Key Using CMC Web Interface	185
Creating Encryption Key Using RACADM	185
Modifying Encryption Key Identifier Using CMC Web Interface	185
Modifying Encryption Identifier Key Using RACADM	185
Deleting Encryption Key Using CMC Web Interface	186
Deleting Encryption Key Using RACADM	186
Encrypting Virtual Disks	186
Encrypting Virtual Disks Using CMC Web Interface	186
Encrypting Virtual Disks Using RACADM	186
Unlocking Foreign Configuration	186
Unlocking Foreign Configuration Using CMC Web Interface	187
Unlocking Foreign Configuration Using RACADM	187
Cryptographic Erase	187
Performing Cryptographic Erase	187
Applying Virtual Adapter Access Policy To Virtual Disks	188
Modifying Virtual Disk Properties Using CMC Web Interface	188
Enclosure Management Module	188
Viewing EMM Status and attributes	189
Viewing Enclosure Status and Attributes	189
Reporting up to two Enclosures per Connector	189
Setting Asset Tag and Asset Name of the Enclosure	190
Viewing Temperature Probe Status and attributes of the Enclosure	190
Setting the Temperature Warning Threshold of the Enclosure	190
Viewing Fan Status and attributes of the Enclosure	191
Viewing Enclosure Properties Using CMC Web Interface	191
Chapter 16: Managing PCIe Slots	192
Viewing PCIe Slot Properties Using CMC Web Interface	192
Assigning PCle Slots To Servers Using CMC Web Interface	193
Managing PCIe Slots Using RACADM	193
PCle Power Ride-Through	194
Viewing PCle Ride-through Properties Using CMC Web Interface	194
Viewing PCle Ridethrough Properties Status Using RACADM	194
Configuring PCle Ride-through Properties Using CMC Web Interface	194
Configuring PCle Ride-through Properties Status Using RACADM	195

Chapter 17: Troubleshooting and Recovery	196
Resetting Forgotten Administrative Password	196
Gathering Configuration Information, Chassis Status, and Logs Using RACDUMP	197
Supported Interfaces	197
Downloading SNMP Management Information Base File	197
First Steps to Troubleshoot a Remote System	198
Power Troubleshooting	198
Troubleshooting Alerts	199
Viewing Event Logs	199
Viewing Hardware Log	199
Viewing Chassis Log	200
Using Diagnostic Console	200
Resetting Components	20′
Saving or Restoring Chassis Configuration	20′
Troubleshooting Network Time Protocol Errors	20′
Interpreting LED Colors and Blinking Patterns	202
Troubleshooting Non-responsive CMC	204
Observing LEDs to Isolate the Problem	204
Obtain Recovery Information from DB-9 Serial Port	204
Recovering Firmware Image	205
Troubleshooting Network Problems	205
Troubleshooting Controller	205
Hotplugging enclosures in fault-tolerant chassis	206
Chapter 18: Using LCD Panel Interface	207
LCD Navigation	
Main Menu	
KVM Mapping Menu	
DVD Mapping	
Enclosure Menu	
IP Summary Menu	
Settings	
Diagnostics	
Front Panel LCD Messages	
LCD Module and Server Status Information	
Chapter 19: Frequently Asked Questions	
RACADM	
Managing and Recovering a Remote System	
Active Directory	
Active DirectoryFlexAddressPlusFlexAddress and FlexAddressPlus	
FlexAddress and FlexAddressPlus	/۱ک ۲۱۵

Overview

The Dell Chassis Management Controller (CMC) for Dell EMC PowerEdge VRTX is a Systems Management hardware and software solution for managing the **PowerEdge VRTX** chassis. The CMC has its own microprocessor and memory and is powered by the modular chassis into which it is plugged.

CMC enables an IT administrator to:

- View inventory
- Perform configuration and monitoring tasks
- Remotely turn on and turn off chassis and servers
- Enable alerts for events on servers and components in the server module
- View and manage storage controller and hard disk drives in the VRTX chassis
- Manage the PCIe subsystem in the VRTX chassis
- Provide a one-many management interface to the iDRACs and I/O modules in the chassis

You can configure the PowerEdge VRTX chassis either with a single CMC, or with redundant CMCs. In redundant CMC configurations, if the primary CMC loses communication with the chassis or the management network, a standby CMC takes over the chassis management.

CMC provides multiple System Management functions for servers. Power and thermal management are the primary functions of CMC, which are listed as follows:

- Enclosure-level real-time automatic power and thermal management.
 - CMC monitors system power requirements and supports the optional Dynamic Power Supply Engagement (DPSE) mode.
 This mode enables CMC to improve power efficiency by setting the power supplies while the server in standby mode and dynamically managing the load and redundancy requirements.
 - o CMC reports real-time power consumption, which includes logging high and low points with a time stamp.
 - o CMC supports setting an optional enclosure maximum power limit (System Input Power Cap), which alerts and takes actions such as limiting the power consumption of servers, and/or preventing the turning on of new servers to keep the enclosure under the defined maximum power limit.
 - CMC monitors and automatically controls the functions of cooling fans and blowers on the basis of actual ambient and internal temperature measurements.
 - o CMC provides comprehensive enclosure inventory and status or error reporting.
- CMC provides a mechanism for centralized configuration of the:
 - o Network and security setting of the Dell PowerEdge VRTX enclosure.
 - o Power redundancy and power ceiling settings.
 - I/O switch and iDRAC network settings.
 - o First boot device on the server module.
 - I/O fabric consistency checks between the I/O module and servers. CMC also disables components, if required, to protect the system hardware.
 - User access security.
 - o Storage components, including the fault-tolerant mode for the storage controllers.
 - o PCle slots.

You can configure CMC to send email alerts or SNMP trap alerts for warnings or errors such as temperature, hardware misconfiguration, power outage, fan speed, and blowers.

Topics:

- What is new in this release
- Key Features
- Chassis Overview
- Minimum CMC Version
- Supported Remote Access Connections
- Supported Platforms
- Supported Web Browsers
- Managing Licenses

- Viewing Localized Versions of the CMC Web Interface
- Supported Management Console Applications
- How to use this Guide
- Other Documents You May Need
- Accessing support content from the Dell EMC support site

What is new in this release

This release of CMC for Dell EMC PowerEdge VRTX supports:

- Displaying fan speed and temperature information using WSMan.
- Integrating the LLDP opensource daemon to forward the LLDP packets to iDRAC through VLAN.
- Transferring CMC dumplogs to iDRAC.
- Unlocking secure foreign configuration using the CMC web interface.
- Displaying wear gauge details of Solid State Drives (SSDs).

Key Features

The CMC features are grouped into management and security features.

Management Features

CMC provides the following management features:

- Redundant CMC environment.
- Dynamic Domain Name System (DDNS) registration for IPv4 and IPv6.
- Login management and configuration for local users, Active Directory, and LDAP.
- Advanced cooling options such and ECM (Enhanced Cooling Mode) and Fan Offset can be enabled to provide additional
 cooling for improved performance.
- Remote system management and monitoring using SNMP, a web interface, KVM, Telnet, or SSH connection.
- Monitoring Provides access to system information and status of components.
- Access to system event logs Provides access to the hardware log and chassis log.
- Firmware updates for various chassis components Enables you to update the firmware for CMC, iDRAC on servers, chassis infrastructure, and chassis storage.
- Firmware update of server components such as BIOS, network controllers, storage controllers, and so on across multiple servers in the chassis using Lifecycle Controller.
- Dell OpenManage software integration Enables you to launch the CMC web interface from Dell OpenManage Server Administrator or OpenManage Essentials (OME) 1.2.
- CMC alert Alerts you about potential managed node issues through Remote syslog email message or SNMP trap.
- Remote power management Provides remote power management functions, such as turn off and reset of any chassis
 component, from a management console.
- Power usage reporting.
- Secure Sockets Layer (SSL) encryption Provides secure remote system management through the web interface.
- Launch point for the Integrated Dell Remote Access Controller (iDRAC) web interface.
- Support for WS-Management.
- FlexAddress feature Replaces the factory-assigned World Wide Name/Media Access Control (WWN/MAC) addresses with chassis-assigned WWN/MAC addresses for a particular slot.
- iDRAC I/O Identity feature support for enhanced WWN/MAC Address Inventory.
- Graphical display of chassis component status and health.
- Support for single and multi-slot servers.
- LCD iDRAC configuration wizard supports for iDRAC network configuration.
- iDRAC single sign-on.
- Network time protocol (NTP) support.
- Enhanced server summary, power reporting, and power control pages.
- Forced CMC failover and virtual reseat of servers.
- Multi-chassis management, allows up to eight other chassis to be visible from the lead chassis.

- Configure storage components on the chassis.
- Map PCle slots to the servers and their identification.

Security Features

The CMC provides the following security features:

- Password-level security management Prevents unauthorized access to a remote system.
- Centralized user authentication through:
 - o Active Directory using Standard Schema or an Extended Schema (optional).
 - o Hardware-stored user IDs and passwords.
- Role-based authority Enables an administrator to configure specific privileges for each user.
- User ID and password configuration through the web interface. Web interface supports 128-bit SSL 3.0 encryption and 40-bit SSL 3.0 encryption (for countries where 128-bit is not acceptable).
 - i NOTE: Telnet does not support SSL encryption.
- Configurable IP ports (if applicable).
- Login failure limits per IP address, with login blocking from the IP address when the limit is exceeded.
- Configurable session auto time out, and more than one simultaneous sessions.
- Limited IP address range for clients connecting to CMC.
- Secure Shell (SSH), which uses an encrypted layer for higher security.
- Single Sign-on, Two-Factor Authentication, and Public Key Authentication.

Chassis Overview

The figure here shows a view of the CMC connectors.




Figure 1. CMC connectors and LEDs

Table 1. CMC connectors and LEDs

Item	Indicator, Button, or Connector
1	Status/identification indicator (CMC 1)
2	Power indicator (CMC 1)
3	CMC connector ports (2)
4	Power indicator (CMC 2)
5	Status/identification indicator (CMC 2)

A Back Panel view of the chassis is given here with a table that lists the parts and devices available in the CMC.




Figure 2. CMC back panel

Table 2. CMC back panel — parts

Item	Indicator, Button, or Connector
1	PCIe expansion card slots low-profile (5)
2	PCIe expansion card slots full height (3)
3	CMC GB Ethernet port (CMC-2)
4	CMC GB Ethernet port (CMC-1)
5	Serial Connector
6	Blower modules (4)
7	I/O module ports
8	PSU 4
9	PSU 3
10	PSU 1
11	PSU 2

A Front Panel view of the chassis is given here with a table that lists the parts and devices available in the CMC.




Figure 3. Front-Panel Features And Indicators—3.5 Inch Hard Disk Drive Chassis

Table 3. Front panel — features and indicators

Item	Indicator, Button, or Connector	Description		
1	USB connectors (2)	Allows a keyboard and mouse to be connected to the system.		
2	LCD panel	Provides system information and status, and error messages to indicate when the system is operating correctly or when the system needs attention.		
3	LCD menu scroll buttons (4)	Moves the cursor in one-step increments.		
4	Selection ("check") button	Selects and saves an item on the LCD screen and moves to the next screen.		
5	Enclosure power-on indicator, power button	The power-on indicator glows when the enclosure power is on. The power button controls the PSU output to the system.		
6	Hard disk drives (HDD)	2.5 inch hard Up to twenty-five 2.5 inch hot-swappable hard disk drives. drive enclosure		
		3.5 inch hard Up to twelve 3.5 inch hot-swappable hard disk drives. drive enclosure		
7	Information tag	A slide-out label panel which allows you to record system information such as Service Tag, NIC, MAC address, the system's electrical rating, and Worldwide Regulatory Agency marks.		
8	Optical drive (optional)	One optional SATA DVD-ROM drive or DVD+/-RW drive.		
9	Vents	Vents for the temperature sensor. NOTE: To make sure about proper cooling, verify that the vents are not blocked.		
10	Video connector	Allows a monitor to be connected to the system.		
11	Server modules	Up to four PowerEdge M520, M620, M630, or M640 server modules or 2 M820 server modules configured for the enclosure.		

Minimum CMC Version

The following table lists the minimum CMC version required to enable the listed server modules.

Table 4. Minimum CMC Version for server modules

Servers	Minimum version of CMC
PowerEdge M520	CMC 1.36
PowerEdge M620	CMC 1.36
PowerEdge M820	CMC 1.36
PowerEdge M630	CMC 2.00
PowerEdge M830	CMC 2.00
PowerEdge M640	CMC 3.00

The following table lists the minimum CMC version required to enable the listed I/O moduless.

Table 5. Minimum CMC Version for I/O modules

IOM Switches	Minimum version of CMC
R1 VRTX 1Gb Pass-through	CMC 1.20
R1-2401 VRTX 1GbE Switch	CMC 1.20
R1-2210 VRTX 10Gb Switch	CMC 2.00

Supported Remote Access Connections

The following table lists the supported Remote Access Controllers.

Table 6. Supported Remote Access Connections

Connection	Features
CMC Network Interface ports	 GB port: Dedicated network interface for the CMC web interface. DHCP support. SNMP traps and e-mail event notification. Network interface for the iDRAC and I/O Modules (IOMs). Support for Telnet/SSH command console and RACADM CLI commands including system boot, reset, power-on, and shutdown commands.
Serial port	 Support for serial console and RACADM CLI commands including system boot, reset, power-on, and shutdown commands. Support for binary interchange for applications designed to communicate with a binary protocol to a particular type of I/O Module. Serial port can be connected internally to the serial console of a server, or I/O module, using the connect (or racadm connect) command. Provides access only to the active CMC.

Supported Platforms

CMC supports modular servers designed for the PowerEdge VRTX platform. For information about compatibility with CMC, see the documentation for your device.

For the latest supported platforms, see the *Dell Chassis Management Controller (CMC) Version 3.0 for Dell PowerEdge VRTX Release Notes* available at **dell.com/support/manuals**.

Supported Web Browsers

The following web browsers are supported for Dell PowerEdge VRTX:

- Microsoft Internet Explorer 9
- Microsoft Internet Explorer 10
- Microsoft Internet Explorer 11
- Microsoft EDGE
- Safari version 7
- Safari version 8
- Safari version 9
- Mozilla Firefox version 52
- Mozilla Firefox version 53
- Google Chrome version 57
- Google Chrome version 58
- NOTE: By default, TLS 1.1 and TLS 1.2 are supported in this release. However, to enable TLS 1.0 use the following racadm command:

\$ racadm config -g cfgRacTuning -o cfgRacTuneTLSProtocolVersionEnable TLSv1.0+

Managing Licenses

The CMC features are available based on the license (CMC Express or CMC Enterprise) purchased. Only licensed features are available in the interfaces that allow you to configure or use CMC. For example, CMC Web interface, RACADM, WS-MAN, and so on. CMC license management and firmware update functionality is always available through CMC Web interface and RACADM.

Types of Licenses

The types of licenses offered are:

- 30 day evaluation and extension The license expires after 30 days that can be extended for 30 days. Evaluation licenses are duration-based, and the timer runs when power is applied to the system.
- Perpetual The license is bound to the service tag and is permanent.

Acquiring Licenses

Use any of the following methods to acquire the licenses:

- E-mail License is attached to an e-mail that is sent after requesting it from the technical support center.
- Self-service portal A link to the Self-Service Portal is available from CMC. Click this link to open the licensing Self-Service
 Portal on the internet from where you can purchase licenses. For more information, see the online help for the self-service
 portal page.
- Point-of-sale License is acquired while placing the order for a system.

License Operations

Before you perform the license management tasks, make sure to acquire the licenses. For more information, see the Overview and Feature Guide available at support.dell.com.

You can perform the following licensing operations using CMC, RACADM, and WS-MAN for one-to-one license management, and Dell License Manager for one-to-many license management:

- i NOTE: If you have purchased a system with all the licenses pre-installed, then license management is not required.
- View View the current license information.
- Import After acquiring the license, store the license in a local storage and import it into CMC using one of the supported interfaces. The license is imported if it passes the validation checks.
 - i NOTE: For a few features, a CMC restart may be required to enable the features.
- Export Export the installed license into an external storage device back up or to reinstall it after a service part is replaced. The file name and format of the exported license is <EntitlementID>.xml
- Delete Delete the license that is assigned to a component if the component is missing. After the license is deleted, it is not stored in CMC and the base product functions are enabled.
- Replace Replace the license to extend an evaluation license, change a license type such as an evaluation license with a purchased license, or extend an expired license.
- An evaluation license may be replaced with an upgraded evaluation license or with a purchased license.
- A purchased license may be replaced with an updated license or with an upgraded license. For more information about license, click Dell Software License Management Portal.
- Learn More Learn more about an installed license, or the licenses available for a component installed in the server.
 - NOTE: For the Learn More option to display the correct page, make sure that *.dell.com is added to the list of Trusted Sites in the Security Settings. For more information, see the Internet Explorer help documentation.

License Component State or Condition and Available Operations

The following table provides the list of license operations available based on the license state or condition.

Table 7. License Operations Based on State and Condition

License/Component state or condition	Import	Export	Delete	Replace	Learn More
Non-administrator login	No	Yes	No	No	Yes
Active license	Yes	Yes	Yes	Yes	Yes
Expired license	No	Yes	Yes	Yes	Yes
License installed but component missing	No	Yes	Yes	No	Yes

Managing Licenses Using CMC Web Interface

To manage the licenses using the CMC Web interface, go to Chassis Overview > Setup > Licenses.

Before importing a license, make sure to store a valid license file on the local system or on a network share that is accessible from the CMC. The license is either embedded, or sent through an email from the **Self-Service Web Portal**, or from the License Key Management tool.

The **Licensing** page displays the licenses that are associated to devices, or the licenses that are installed, but the device is not present in the system. For more information about importing, exporting, deleting, or replacing a license, see the *Online Help*.

Managing Licenses Using RACADM

To manage licenses using the RACADM commands, use the following license subcommand.

racadm license < license command type>

For more information about the RACADM commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at dell.com/support/Manuals.

Licensable Features in CMC

A list of CMC features that are enabled on the basis of your license is given here in the table.

Table 8. Licensable features

Feature	Express	Enterprise	Notes
CMC Network	Yes	Yes	
CMC Serial Port	Yes	Yes	
RACADM (SSH, Local, and Remote)	Yes	Yes	
CMC Setup Backup	No	Yes	
CMC Setup Restore	Yes	Yes	
WS-MAN	Yes	Yes	
SNMP	Yes	Yes	
Telnet	Yes	Yes	
SSH	Yes	Yes	
Web-based Interface	Yes	Yes	
Email Alerts	Yes	Yes	
LCD Deployment	Yes	Yes	
Extended iDRAC Management	Yes	Yes	
Remote Syslog	No	Yes	
Directory Services	No*	Yes	*For non-default directory service setting, only Reset Directory Services is allowed with an Express license. Reset Directory Services will set the Directory services to factory default.
iDRAC Single Sign-On	No	Yes	
Two-Factor Authentication	No	Yes	
PK Authentication	No	Yes	
Remote File Share	Yes	Yes	
Slot Resource Management	No	Yes	
Enclosure-level power capping	No*	Yes	*For non-default power cap setting, only Restore Power Cap is allowed with an Express license. Restore Power Cap will reset the Power Cap settings to factory default.
Dynamic Power Supply Engagement	No*	Yes	*For non-default DPSE settings, only Restore DPSE is allowed with an Express license. Restore DPSE will reset the DPSE to factory default.
Multi-chassis management	No	Yes	

Table 8. Licensable features (continued)

Feature	Express	Enterprise	Notes
Advanced Configuration	No	Yes	
Enclosure-level backup	No	Yes	
FlexAddress Enablement	No*	Yes	*For non-default FlexAddress settings, only Restore Default is allowed with Express license. Restore Default will reset the FlexAddress settings to the factory default.
PCle Adapter Mapping	Yes*	Yes	*A maximum of two PCle Adapters can be assigned per Server with Express License.
Virtual Adapter to Slot Mapping	No*	Yes	*For non-default mapping of Virtual Adapters, only Default mapping is allowed with an Express license. Restore Default will change virtual adapter mapping to factory default.
Virtual Adapter to Slot UnMapping	Yes	Yes	
Server cloning	No	Yes	
One-to-many Server Firware Update	No	Yes	
One-to-many configuration for iDRAC	No	Yes	
Boot Identity	No	Yes	
Chassis Profile	No	Yes	
Quick Deploy	No	Yes	

Viewing Localized Versions of the CMC Web Interface

To view localized versions of the CMC web interface, read through your web browser's documentations.

Supported Management Console Applications

CMC supports integration with Dell OpenManage Console. For more information, see the OpenManage Console documentation available at **dell.com/support/manuals**.

How to use this Guide

The contents of this User's Guide enable you to perform the tasks by using:

- The Web interface: Only the task-related information is given here. For information about the fields and options, see the CMC for Dell PowerEdge VRTX Online Help that you can open from the Web interface.
- The RACADM commands: The RACADM command or the object that you must use is provided here. For more information
 about a RACADM command, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line
 Reference Guide available at dell.com/cmcmanuals.

Other Documents You May Need

To access the documents from the Dell Support site. Along with this Reference Guide, you can access the following guides available at **dell.com/support/manuals**.

- The VRTX CMC Online Help provides information about using the Web interface. To access the Online Help, click **Help** on the CMC web interface.
- The Chassis Management Controller Version 3.0 for Dell PowerEdge VRTX RACADM Command Line Reference Guide provides information about using the VRTX-related RACADM features.
- The Dell Chassis Management Controller (CMC) for Dell PowerEdge VRTX Version 3.0 Release Notes, available at **dell.com/cmcmanuals**, provides last-minute updates to the system or documentation or advanced technical reference material intended for experienced users or technicians.
- The Integrated Dell Remote Access Controller (iDRAC) User's Guide provides information about installation, configuration, and maintenance of the iDRAC on managed systems.
- The Dell OpenManage Server Administrator's User's Guide provides information about installing and using Server Administrator.
- The Dell OpenManage SNMP Reference Guide for iDRAC and Chassis Management Controller provides information about SNMP MIBs.
- The Dell Update Packages User's Guide provides information about obtaining and using Dell Update Packages as part of your system update strategy.
- The Dell Shared PowerEdge RAID Controller (PERC) 8 User's Guide provides information about deploying the Shared PERC 8 card and managing the storage subsystem. This document is available online at dell.com/storagecontrollermanuals.
- Dell systems management application documentation provides information about installing and using the systems management software.

The following system documents provide more information about the system in which VRTX CMC is installed:

- The safety instructions that came with your system provide important safety and regulatory information. For additional regulatory information, see the Regulatory Compliance home page at www.dell.com/regulatory_compliance. Warranty information may be included within this document or as a separate document.
- The Dell PowerEdge VRTX Getting Started Guide shipped with your system provides an overview of system features, setting up your system, and technical specifications.
- The setup placemat shipped with your system provides information about the initial system setup and configuration.
- The server module's *Owner's Manual* provides information about the server module's features and describes how to troubleshoot the server module and install or replace the server module's components. This document is available online at **dell.com/poweredgemanuals**.
- The rack documentation included with your rack solution describes how to install your system into a rack, if required.
- For the full name of an abbreviation or acronym used in this document, see the Glossary at dell.com/support/manuals.
- Systems management software documentation describes the features, requirements, installation, and basic operation of the software.
- Documentation for any components you purchased separately provides information to configure and install these options.
- Any media that ships with your system that provides documentation and tools for configuring and managing your system, including those pertaining to the operating system, system management software, system updates, and system components that you purchased with your system. For more information on the system, scan the Quick Resource Locator (QRL) available on your system and the system setup placemat that shipped with your system. Download the QRL application from your mobile platform to enable the application on your mobile device.

Accessing support content from the Dell EMC support site

Access supporting content related to an array of systems management tools using direct links, going to the Dell EMC support site, or using a search engine.

- Direct links:
 - For Dell EMC Enterprise Systems Management and Dell EMC Remote Enterprise Systems Management—https://www.dell.com/esmmanuals
 - For Dell EMC Virtualization Solutions—https://www.dell.com/SoftwareManuals
 - o For Dell EMC OpenManage—https://www.dell.com/openmanagemanuals
 - For iDRAC—https://www.dell.com/idracmanuals

- For Dell EMC OpenManage Connections Enterprise Systems Management—https://www.dell.com/ OMConnectionsEnterpriseSystemsManagement
- o For Dell EMC Serviceability Tools—https://www.dell.com/serviceabilitytools
- Dell EMC support site:
 - 1. Go to https://www.dell.com/support.
 - 2. Click Browse all products.
 - 3. From the All products page, click Software, and then click the required link.
 - **4.** Click the required product and then click the required version.

Using search engines, type the name and version of the document in the search box.

Installing and Setting Up CMC

This section provides information about how to install your CMC hardware, establish access to CMC, configure your management environment to use CMC, and guides you through the tasks for configuring a CMC:

- Set up initial access to CMC.
- Access CMC through a network.
- Add and configure CMC users.
- Update CMC firmware.

For more information about installing and setting up redundant CMC environments, see Understanding Redundant CMC Environment.

Topics:

- Before You Begin
- Installing CMC Hardware
- Installing Remote Access Software on a Management Station
- Configuring a Web Browser
- Setting Up Initial Access to CMC
- Interfaces and Protocols to Access CMC
- Downloading and Updating CMC Firmware
- Setting Chassis Physical Location and Chassis Name
- Setting Date and Time on CMC
- Configuring LEDs to Identify Components on the Chassis
- Configuring CMC Properties
- Configuring iDRAC Launch Method Using CMC Web Interface
- Configuring iDRAC Launch Method Using RACADM
- Configuring Login Lockout Policy Attributes Using CMC Web Interface
- Configuring Login Lockout Policy Attributes Using RACADM
- Understanding Redundant CMC Environment
- Configuring Front Panel

Before You Begin

Before setting up your CMC environment, download the latest version of CMC firmware for PowerEdge VRTX from **dell.com/support/**.

Also, make sure that you have the *Dell Systems Management Tools and Documentation DVD* that was included with your system.

Installing CMC Hardware

CMC is pre-installed on your chassis and hence no installation is required. You can install a second CMC to run as a standby to the active CMC.

Checklist To Set up Chassis

The following tasks enable you to accurately set up the chassis:

- 1. CMC and the management station, where you use your browser, must be on the same network, which is called the management network. Connect an Ethernet network cable from the CMC active port to the management network.
- 2. Install the I/O module in the chassis and connect the network cable to the chassis.

- 3. Insert the servers in the chassis.
- 4. Connect the chassis to the power source.
- 5. Press the power button, or turn on the chassis from the CMC web interface after completing the task in step 7.
 - i NOTE: Do not turn on the servers.
- 6. Using the LCD panel, navigate to the IP Summary and click on the Check button to select. Use the IP address for the CMC in the management system browser (IE, Chrome, or Mozilla). To set up DHCP for CMC, use LCD panel to, click Main Menu > Settings > Network Settings.
- 7. Connect to the CMC IP address by using a web browser by typing the default username (root) and password (calvin).
- 8. Provide each iDRAC with an IP address in the CMC web interface, and enable the LAN and IPMI interface.
 - NOTE: iDRAC LAN interface on some servers are disabled by default. This information can be found on the CMC web interface under **Server Overview** > **Setup**. This might be an advanced license option; in which case you must use the **SetUp** feature for each server).
- 9. Provide the IO module with an IP address in the CMC web interface. You can get the IP address by clicking I/O Module Overview, and then clicking Setup.
- 10. Connect to each iDRAC through the Web browser and provide final configuration of iDRAC. The default user name is root and password is calvin.
- 11. Connect to the I/O module by using the web browser and provide final configuration of the IO module.
- 12. Turn on the servers and install the operating system.

Basic CMC Network Connection

For the highest degree of redundancy, connect each available CMC to your management network.

Installing Remote Access Software on a Management Station

You can access CMC from a management station using remote access software, such as Telnet, Secure Shell (SSH), or serial console utilities provided on your operating system or using the web interface.

To use remote RACADM from your management station, install remote RACADM using the *Dell Systems Management Tools and Documentation* DVD that is available with your system. This DVD includes the following Dell OpenManage components:

- DVD root Contains the Dell Systems Build and Update Utility.
- SYSMGMT Contains the systems management software products including Dell OpenManage Server Administrator.
- Docs Contains documentation for systems, systems management software products, peripherals, and RAID controllers.
- SERVICE Contains the tools required to configure your system, and delivers the latest diagnostics and Dell-optimized drivers for your system.

For information about installing Dell OpenManage software components, see the *Dell OpenManage Installation and Security User's Guide* available on the DVD or at **dell.com/support/manuals**. You can also download the latest version of the Dell DRAC Tools from **support.dell.com**.

Installing RACADM on a Linux Management Station

- 1. Log in as root to the system running a supported Red Hat Enterprise Linux or SUSE Linux Enterprise Server operating system where you want to install the managed system components.
- 2. Insert the Dell Systems Management Tools and Documentation DVD into the DVD drive.
- 3. To mount the DVD to a required location, use the mount command or a similar command.
 - NOTE: On the Red Hat Enterprise Linux 5 operating system, DVDs are auto-mounted with the -noexec mount option. This option does not allow you to run any executable from the DVD. You need to mount the DVD-ROM manually, and then run the commands.
- Navigate to the SYSMGMT/ManagementStation/linux/rac directory. To install the RAC software, type the following command:

```
rpm -ivh *.rpm
```

- For help about the RACADM command, type racadm help after you run the previous commands. For more information
 about RACADM, see the Chassis Management Controller for Dell PowerEdge VRTX RACADM Command Line Reference
 Guide.
 - NOTE: When using the RACADM remote capability, you must have the 'write' permission on the folders where you are using the RACADM subcommands, involving the file operations. For example, racadm getconfig -f <file name>.

Uninstalling RACADM From a Linux Management Station

- 1. Log in as root to the system where you want to uninstall the management station features.
- 2. Run the following rpm query command to determine which version of the DRAC tools is installed: rpm -qa | grep mgmtst-racadm
- 3. Verify the package version to be uninstalled and uninstall the feature by using the rpm -e rpm -qa | grep mgmtst-racadm command.

Configuring a Web Browser

You can configure and manage CMC, servers, and modules installed in the chassis through a web browser. See the "Supported Browsers" section in the *Dell Systems Software Support Matrix* at **dell.com/support/manuals**.

The CMC and the management station where you use your browser must be on the same network, which is called the *management network*. On the basis of your security requirements, the management network can be an isolated and highly secure network.

NOTE: Make sure that the security measures on the management network such as firewalls and proxy servers, do not prevent your web browser from accessing the CMC.

Some browser features can interfere with connectivity or performance, especially if the management network does not have a route to the Internet. If your management station is running on a Windows operating system, some Internet Explorer settings can interfere with connectivity, even though you use a command line interface to access the management network.

NOTE: To address security issues, Microsoft Internet Explorer strictly monitors the time on its cookie management. To support this, the time on your computer that runs Internet Explorer must be synchronized with the time on the CMC.

Proxy Server

To browse through a proxy server that does not have access to the management network, you can add the management network addresses to the exception list of the browser. This instructs the browser to bypass the proxy server while accessing the management network.

Internet Explorer

To edit the exception list in Internet Explorer:

- 1. Start Internet Explorer.
- 2. Click Tools > Internet Options > Connections.
- 3. In the Local Area Network (LAN) settings section, click LAN Settings.
- 4. In the Proxy server section, select the Use a proxy server for your LAN (These settings will not apply to dial-up or VPN connections) option, and then click Advanced.
- **5.** In the **Exceptions** section, add the addresses for CMCs and iDRACs on the management network to the semicolon-separated list. You can use DNS names and wildcards in your entries.

Mozilla FireFox

To edit the exception list in Mozilla Firefox version 19.0:

- 1. Start Mozilla Firefox.
- 2. Click Tools > Options (for systems running on Windows), or click Edit > Preferences (for systems running on Linux).
- 3. Click Advanced, and then click the Network tab.
- 4. Click Settings.
- 5. Select Manual Proxy Configuration.
- 6. In the **No Proxy for** field, type the addresses for CMCs and iDRACs on the management network to the comma-separated list. You can use DNS names and wildcards in your entries.

Microsoft Phishing Filter

If the Microsoft Phishing Filter is enabled in Internet Explorer on your management system, and your CMC does not have Internet access, accessing CMC may be delayed by a few seconds. This delay can happen if you are using the browser or another interface such as remote RACADM. To disable the phishing filter:

- 1. Start Internet Explorer.
- 2. Click Tools > Phishing Filter, and then click Phishing Filter Settings.
- 3. Select the Disable Phishing Filter option and click OK.

Certificate Revocation List Fetching

If your CMC has no access to the Internet, disable the certificate revocation list (CRL) fetching feature in Internet Explorer. This feature tests whether a server such as the CMC web server uses a certificate that is on a list of revoked certificates retrieved from the Internet. If the Internet is inaccessible, this feature can cause delays of several seconds when you access the CMC using the browser or with a command line interface such as remote RACADM.

To disable CRL fetching:

- 1. Start Internet Explorer.
- 2. Click Tools > Internet Options, and then click Advanced.
- 3. Go to the Security section, clear the Check for publisher's certificate revocation option, and then click OK.

Downloading Files From CMC With Internet Explorer

When you use Internet Explorer to download files from the CMC, you may experience problems when the **Do not save encrypted pages to disk** option is not enabled.

To enable the **Do not save encrypted pages to disk** option:

- 1. Start Internet Explorer.
- 2. Click Tools > Internet Options > Advanced.
- 3. In the Security section, select the Do not save encrypted pages to disk option.

Enabling Animations In Internet Explorer

When transferring files to and from the web interface, a file transfer icon spins to show transfer activity. While using Internet explorer, you have to configure the browser to play animations.

To configure Internet Explorer to play animations:

- 1. Start Internet Explorer.
- 2. Click Tools > Internet Options > Advanced.
- 3. Go to the Multimedia section, and then select the Play animations in web pages option.

Setting Up Initial Access to CMC

To remotely manage the CMC, connect the CMC to your management network, and then configure the CMC network settings.

(i) NOTE: To manage the PowerEdge VRTX solution, it must be connected to your management network.

For information about configuring CMC network settings, see Configuring Initial CMC Network. This initial configuration assigns the TCP/IP networking parameters that enable access to CMC.

CMC and iDRAC on each server and the network management ports for the switch I/O module are connected to a common integrated network in the PowerEdge VRTX chassis. This allows the management network to be isolated from the server data network. It is important to separate this traffic for uninterrupted access to chassis management.

CMC is connected to the management network. All external access to CMC and iDRACs is achieved through CMC. Access to the managed servers, conversely, is accomplished through network connections to the I/O module (IOM). This allows the application network to be isolated from the management network.

It is recommended to isolate chassis management from the data network. Due to the potential of traffic on the data network, the management interfaces on the internal management network can be saturated by traffic intended for servers. This results in CMC and iDRAC communication delays. These delays may cause unpredictable chassis behavior, such as CMC displaying iDRAC as offline even when it is up and running, which in turn causes other unwanted behavior. If physically isolating the management network is impractical, the other option is to separate CMC and iDRAC traffic to a separate VLAN. CMC and individual iDRAC network interfaces can be configured to use a VLAN.

Configuring Initial CMC Network

i NOTE: Changing your CMC Network settings may disconnect your current network connection.

You can perform the initial network configuration of CMC before or after CMC has an IP address. If you configure CMC's initial network settings before you have an IP address, you can use either of the following interfaces:

- The LCD panel on the front of the chassis
- Dell CMC serial console

If you configure initial network settings after the CMC has an IP address, you can use any of the following interfaces:

- Command line interfaces (CLIs) such as a serial console, Telnet, SSH, or the Dell CMC console.
- Remote RACADM
- CMC web interface
- LCD Panel interface

CMC supports both IPv4 and IPv6 addressing modes. The configuration settings for IPv4 and IPv6 are independent of each other.

Configuring CMC Network Using LCD Panel Interface

You can use the LCD panel interface to set up the CMC network.

- (i) NOTE: You can customize the orientation of an LCD display (for rack or tower mode) by keeping the up-down buttons pressed for two seconds. Alternately, you can also use the right-left buttons. For more information about the buttons available on a CMC LCD panel, see LCD Navigation.
- 1. To start the CMC configuration:
 - For a chassis that has not been configured earlier, the **LCD Language** panel is displayed. On the **LCD Language** panel, navigate to the required language using the arrow buttons. When the desired language is highlighted, select the language by pressing the Center button. The **Network Settings** panel is displayed.
 - For a chassis that has been configured earlier, the **Main Menu** panel is displayed. From the **Main Menu**, select **Settings** and then **Network Settings**.
- 2. On the **Network Settings** panel, select the required mode of setup:
 - **Quick Setup (DHCP)** Select this mode to set up CMC quickly using DHCP addresses. For information on configuring CMC using this mode see **Configuring CMC Using Quick Setup (DHCP)**.
 - Advanced Setup Select this mode to set up CMC for advanced configurations. For information on configuring CMC using this mode see Configuring CMC Using Advanced Setup.

Configuring CMC Using Quick Setup (DHCP)

To set up a network using the LCD panel interface:

1. From the Network Settings panel, select Quick Setup (DHCP). The panel displays the following message.

About to get DHCP addresses. Ensure CMC network cable is connected.

2. Press the center button to highlight the accept button. Press the center button again to accept the settings or navigate to the back arrow and the press the center button to go back and modify the settings.

Configuring CMC Using Advanced Setup

1. On the **Network Settings** panel, if you select **Advanced Setup**, the following message is displayed to confirm if want to configure CMC:

Configure CMC?

- 2. To configure CMC by using advanced setup properties, click the center button selecting the check icon.
 - NOTE: To skip the CMC configuration navigate to the 'X' icon and then press the center button.
- 3. If you are asked to select an appropriate network speed, select a network speed (Auto (1Gb), 10Mb, or 100Mb) using appropriate buttons.

For effective network throughput, the network speed setting must match your network configuration. Setting the network speed lower than the speed of your network configuration increases bandwidth consumption and slows down the network communication. Determine whether or not your network supports the above network speeds and set it accordingly. If the network configuration does not match any of these values, it is recommended to select the **Auto (1 Gb)** option, or refer to your network equipment manufacturer's user documentation.

- 4. Perform one of the following tasks:
 - Select **Auto (1Gb)**, by pressing the center button, and then pressing the center button again. The **Protocol** panel is displayed. Go to step 6.
 - Select 10Mb, or 100Mb. The Duplex panel

is displayed. Go to step 5.

Else, if you

- 5. On the **Duplex** panel, to select the duplex mode (**Full** or **Half**) that matches the network environment, press the center button, and then press the center button. The **Protocol** panel is displayed.
 - (1Gbps) is selected. If auto negotiation is turned on for one device but not the other, the device that is using auto negotiation can determine the network speed of the other device, but not the duplex mode. In this case, half duplex is selected as the duplex mode during auto negotiation. Such a duplex mismatch results in a slow network connection.
- 6. On the **Protocol** panel, select an Internet Protocol (**IPv4 Only**, **IPv6 Only**, or **Both**) that you want to use for CMC, press the center button, and then press the center button.
- 7. If you select **IPv4** or **Both**, select **DHCP** or **Static** mode. Go to step 8.
 - Else, if you select IPv6, the Configure iDRAC panel is displayed. Go to step 11 later in this procedure.
- 8. On the **Mode** panel, select the mode in which the CMC must obtain the NIC IP addresses. If you select DHCP, CMC retrieves IP configuration (IP address, mask, and gateway) automatically from a DHCP server on your network. CMC is assigned a unique IP address that is allotted on your network. If you select **DHCP**, press the center button, and then press the center again. The **Configure iDRAC** panel is displayed. Go to step 11 later in this procedure.
- 9. If you select **Static**, enter the IP address, gateway, and subnet mask by following the instructions on the LCD panel.

The IP information that you entered is displayed. Press the center button, and then press the center button again. The **CMC Configuration** screen lists the **Static IP Address**, **Subnet Mask**, and **Gateway** settings you entered. Check the settings for accuracy. To correct a setting, press appropriate buttons. Press the center button, and then press the center button. The **Register DNS?** panel is displayed.

- 10. To register select the check icon and press the center button. Set the DNS IP address, select the check icon, and then press the center button. If DNS registration is not required, then select the 'X' icon and press the center button.
- 11. Indicate whether or not you want to configure iDRAC:
 - No: Select the 'X' icon then press the center button. Go to step 17 later in this procedure.
 - Yes: Select the check icon then press the center button.

You can also configure iDRAC from the CMC web interface.

- 12. On the Protocol panel, select the IP type that you want to use for the servers:
 - IPv4 The options DHCP or Static are displayed.
 - Both
 - The options DHCP or Static are displayed.
 - IPv6
 - The **iDRAC Configuration** panel is displayed. Go to step 15.
- 13. Select DHCP or Static.

Table 9. Network mode

Dynamic Host Configuration Protocol (DHCP)	iDRAC retrieves IP configuration (IP address, mask, and gateway) automatically from a DHCP server on your network. The iDRAC is assigned a unique IP address allotted over your network. Press the center button. The IPMI Over LAN panel is displayed.			
Static	If you select Static , manually enter the IP address, gateway, and subnet mask by following the instructions on the LCD screen. If you have selected the Static option, press the center button, and then do the following:			
	a. The following message asks you whether or not you want to automatically increment by using the IP of Slot-1. IPs will auto-increment by slot number.			
	Click the center button. The following message asks you to enter the slot-1 IP number. Enter slot 1 (starting) IP			
	Enter the slot-1 IP number, and then press the center button. b. Set the Subnet mask, and then press the center button. c. Set the gateway, and then press the center button. d. The Network Summary screen lists the Static IP Address, Subnet Mask, and Gateway settings you entered. Check the settings for accuracy. To correct a setting, press appropriate buttons, and then press the center button. e. When you have confirmed the accuracy of the settings you entered, go to step 10. The IPMI Over LAN panel is displayed.			

- 14. From the The IPMI Over LAN panel, select **Enable** or **Disable** to enable or disable IPMI over LAN. Press the center button to continue.
- 15. On the iDRAC Configuration panel, the following message is displayed.

```
Apply settings to installed servers?
```

To apply all iDRAC network settings to the installed servers, select the check icon, and then press the center button. Else, select the 'X' icon and press the center button.

16. On the next iDRAC Configuration panel, the following message is displayed.

```
Auto-Apply settings to newly-inserted servers?
```

To apply all iDRAC network settings to the newly installed servers, select the check icon and press the center button. When a new server is inserted in the chassis, the LCD prompts you whether or not to automatically deploy the server using the previously-configured network settings policies. If you do not want to apply the iDRAC network settings to newly-installed servers, select the 'X' icon and press the center button. When a new server is inserted in the chassis, the iDRAC network settings do not get configured.

17. On the iDRAC Configuration panel, the following message is displayed.

```
Apply All Enclosure Settings?
```

To apply all enclosure settings, select the check icon and press the center button. Else, select the 'X' icon and press the center button.

18. On the **IP Summary** panel, after the 30 second wait panel, review the IP addresses you provided to make sure the addresses are correct. To correct a setting, press the left arrow icon, and then press the center key to return to the screen for that setting. After correcting an IP address, press the center button.

When you have confirmed that the settings you entered are accurate, press the center button, and then press the center button. The **Main Menu** panel id displayed.

CMC and iDRACs are now available on the network. You can access the CMC on the assigned IP address using the Web interface or CLIs such as a serial console, Telnet, and SSH.

Interfaces and Protocols to Access CMC

After you have configured the CMC network settings, you can remotely access CMC using various interfaces. The following table lists the interfaces that you can use to remotely access CMC.

- NOTE: Telnet is not as secure as the other interfaces, hence it is disabled by default. Enable Telnet by using web, SSH, or remote RACADM.
- (i) NOTE: Using more than one interface at the same time may generate unexpected results.

Table 10. CMC Interfaces

Interface	Description		
Web interface	Provides remote access to CMC using a graphical user interface. The Web interface is built into the CMC firmware and is accessed through the NIC interface from a supported web browser on the management station.		
	For a list of supported Web browsers, see the "Supported Browsers" section in the Dell System Software Support Matrix at dell.com/support/manuals.		
Remote RACADM command line interface	 Use this command line utility to manage CMC and its components. You can use remote or firmware RACADM: Remote RACADM is a client utility that runs on a management station. It uses the out-of-band network interface to run RACADM commands on the managed system and uses the HTTPs channel. The -r option runs the RACADM command over a network. Firmware RACADM is accessible by logging in to CMC using SSH or Telnet. You can run the firmware RACADM commands without specifying the CMC IP, user name, or password. After you enter the RACADM prompt, you can directly run the commands without the racadm prefix. 		
Chassis LCD Panel	Use the LCD on the front panel to: View alerts and CMC IP. Set DHCP. Configure CMC static IP settings. View CMC MAC address for the active CMC. View the CMC VLAN ID appended to the end of CMC IP, if the VLAN is already configured.		
Telnet	Provides command line access to CMC through the network. The RACADM command line interface and the connect command, which is used to connect to the serial console of a server or IO module, are available from the CMC command line. (i) NOTE: Telnet is not a secure protocol and is disabled by default. Telnet transmits all data, including passwords in plain text. When transmitting sensitive information, use the SSH interface.		
SSH	Use SSH to run RACADM commands. It provides the same capabilities as the Telnet console using an encrypted transport layer for higher security. The SSH service is enabled by default on CMC and can be disabled.		
WSMan	The WSMan Services are based on the Web Services for Management (WSMan) protocol to perform one-to-many systems management tasks. You must use WSMan client such as WinRM client (Windows) or the OpenWSMan client (Linux) to use the LC-Remote Services functionality. You can also use Power Shell and Python script the WSMan interface.		

Table 10. CMC Interfaces (continued)

Interface	Description
	WSMan is a Simple Object Access Protocol (SOAP)—based protocol used for systems management. CMC uses WS—Management to convey Distributed Management Task Force (DMTF) Common Information Model (CIM)—based management information. The CIM information defines the semantics and information types that can be modified in a managed system.
	The CMC WSMan implementation uses SSL on port 443 for transport security, and supports basic authentication. The data available through WS-Management is provided by CMC instrumentation interface mapped to the DMTF profiles and extension profiles.
	For more information, see: MOFs and Profiles — delltechcenter.com/page/DCIM.Library DTMF Web site — dmtf.org/standards/profiles/ WSMan Release notes file. www.wbemsolutions.com/ws_management.html
	DMTF WSManagement Specifications: www.dmtf.org/standards/wbem/wsman
	Web services interfaces can be utilized by leveraging client infrastructure, such as Windows WinRM and Powershell CLI, open source utilities like WSManCLI, and application programming environments like Microsoft .NET.
	For client connection using Microsoft WinRM, the minimum required version is 2.0. For more information, refer to the Microsoft article, <support.microsoft.com 968929="" kb="">.</support.microsoft.com>

(i) NOTE: The default values of CMC user name and password are root and calvin respectively.

Launching CMC Using Other Systems Management Tools

You can also launch CMC from the Dell Server Administrator or Dell OpenManage Essentials.

To access CMC interface using Dell Server Administrator, launch Server Administrator on your management station. In the left pane of the Server Administrator home page, click **System > Main System Chassis > Remote Access Controller**. For more information, see the *Dell Server Administrator User's Guide* at **dell.com/support/manuals**.

Downloading and Updating CMC Firmware

To download the CMC firmware, see Downloading CMC Firmware.

To update the CMC firmware, see Updating CMC Firmware.

Setting Chassis Physical Location and Chassis Name

You can set the chassis location in a data center and the chassis name to identify the chassis on the network (default name is **Dell Rack System**). For example, an SNMP query on the chassis name returns the name you configure.

Setting Chassis Physical Location and Chassis Name Using Web Interface

To set the chassis location and chassis name using the CMC web interface:

- 1. In the left pane, go to Chassis Overview, and then click Setup .
- 2. On the **General Chassis Settings** page, type the location properties and the chassis name. For more information about setting chassis properties, see the *CMC Online Help*.

- NOTE: The Chassis Location field is optional. It is recommended to use the Data Center, Aisle, Rack, and Rack Slot fields to indicate the physical location of the chassis.
- 3. Click **Apply**. The settings are saved.

Setting Chassis Physical Location and Chassis Name Using RACADM

To set the chassis name, location, date, and time by using the command line interface, see the **setsysinfo** and **setchassisname** commands. For more information, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.*

Setting Date and Time on CMC

You can manually set the date and time, or you can synchronize the date and time with a Network Time Protocol (NTP) server.

Setting Date and Time on CMC Using CMC Web Interface

To set the date and time on CMC:

- 1. In the left pane, click Chassis Overview > Setup > Date/Time.
- To synchronize the date and time with a Network Time Protocol (NTP) server, on the Date/Time page, select Enable NTP
 and specify up to three NTP servers. To manually set the date and time, clear the Enable NTP option, and then edit the
 Date and Time fields.
- 3. Select the **Time Zone** from the drop-down menu, and then click **Apply**.

Setting Date and Time on CMC Using RACADM

To set the date and time using the command line interface, see the **config** command and cfgRemoteHosts database property group sections in the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/support/manuals**.

Configuring LEDs to Identify Components on the Chassis

You can enable the LEDs of components (chassis, servers, physical disk drives, virtual disks, and I/O Modules) to blink so that you can identify the component on the chassis.

i NOTE: To modify these settings, you must have the Chassis Configuration Administrator privilege.

Configuring LED Blinking Using CMC Web Interface

To enable blinking for one, multiple, or all component LEDs:

- In the left pane, go to any of the following pages:
 - o Chassis Overview > Troubleshooting.
 - o Chassis Overview > Chassis Controller > Troubleshooting.
 - Chassis Overview > Server Overview > Troubleshooting.
 - i NOTE: Only servers can be selected on this page.
 - o Chassis Overview > I/O Module Overview > Troubleshooting.
 - Storage > Troubleshooting > Identify.

NOTE: Physical disk per enclosures, virtual disks per enclosures, and external storage component LED can be selected on this page.

To enable blinking of a component LED, select the **Select/Deselect All** option corresponding to the physical disk drive or virtual disk or enclosures, and then click **Blink**. To disable blinking of a component LED, clear the **Select/Deselect All** option corresponding to the LED, and then click **Unblink**.

Configuring LED Blinking Using RACADM

Open a serial/Telnet/SSH text console to CMC, log in, and type:

racadm setled -m < module > [-l < ledState >], where < module > specifies the module whose LED you want to configure. Configuration options:

- server-n where n = 1-4
- switch-1
- cmc-active

and <ledState> specifies whether or not the LED should blink. Configuration options:

- 0 not blinking (default)
- 1 blinking

racadm raid component FQDD>, where the operation value is blink or unblink, and the FQDD is for the component's physical disk drive, virtual disk and enclosures.

Configuring CMC Properties

You can configure CMC properties such as power budgeting, network settings, users, and SNMP and email alerts using the web interface or RACADM commands.

Configuring iDRAC Launch Method Using CMC Web Interface

To configure the iDRAC launch method from the **General Chassis Settings** page:

- In the left pane, click Chassis Overview > Setup.
 The General Chassis Settings page is displayed.
- 2. In the drop-down menu for the iDRAC Launch Method property, select IP Address or DNS.
- 3. Click Apply
 - NOTE: A DNS-based launch is used for any particular iDRAC only if:
 - The chassis setting is DNS.
 - CMC has detected that the specific iDRAC is configured with a DNS name.

Configuring iDRAC Launch Method Using RACADM

To update CMC firmware using RACADM, use the cfgRacTuneIdracDNSLaunchEnable subcommand. For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Configuring Login Lockout Policy Attributes Using CMC Web Interface

i NOTE: To perform the following tasks, you must have Chassis Configuration Administrator privilege.

The **Log in Security** enables you to configure the IP range attributes for CMC login using the CMC web interface. To configure the IP range attributes using CMC web interface:

- In the left pane, go to Chassis Overview and click Network > Network.
 The Network Configuration page is displayed.
- In the IPv4 Settings section, click Advanced Settings. Alternatively, to access the Log in Security page, in the left pane, go to Chassis Overview, click Security > Log in.
 The Log in Security page is displayed.
- 3. To enable the user blocking or IP blocking feature, in the Login Lockout Policy section, select Lockout by User Name or Lockout by IP Address (IPV4).
 - The options to set the other login lockout policy attributes are activated.
- 4. Enter the required values for login lockout policy attributes in the activated fields Lockout Fail Count, Lockout Fail Window, and Lockout Penalty Time. For more information, see the CMC Online Help.
- 5. To save these settings, click **Apply**.

Configuring Login Lockout Policy Attributes Using RACADM

You can use RACADM to configure the Login lockout policy attributes for the following features:

- User blocking
- IP address blocking
- Number of login attempts allowed
- Timespan for the lockout failure counts to occur
- Lockout penalty time
- To enable user blocking feature, use:

```
racadm config -g cfgRacTuning -o cfgRacTuneUserBlkEnable <0|1>
```

• To enable IP blocking feature, use:

```
racadm config -g cfgRacTuning -o cfgRacTuneIPBlkEnable <0|1>
```

• To specify the number of login attempts, use:

```
racadm config -g cfgRacTuning -o cfgRacTuneIpBlkFailCount
```

• To specify the time span within which, lockout fail count failures must occur, use:

```
racadm config -g cfgRacTuning -o cfgRacTuneIpBlkFailWindow
```

• To specify value for lockout penalty time, use:

```
racadm config -g cfgRacTuning -o cfgRacTuneIpBlkPenaltyTime
```

For more information about these objects, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Understanding Redundant CMC Environment

You can install a standby CMC that takes over if your active CMC stops functioning. The redundant CMC may be pre-installed or can be installed later. To make sure full redundancy or best performance, it is important that the CMC network is properly cabled.

Failovers can occur when you:

- Run the RACADM cmcchangeover command. See the cmcchangeover command section in the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.
- Run the RACADM racreset command on the active CMC. See the racreset command section in the Chassis
 Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/
 manuals..

- Reset the active CMC from web interface. See the Reset CMC option for Power Control Operations that is described in Executing Power Control Operations.
- Remove the network cable from the active CMC.
- Remove the active CMC from the chassis.
- Initiate a CMC firmware flash on the active CMC.
- Have an active CMC that is no longer functional.
 - NOTE: In the event of a CMC failover, all iDRAC connections and all active CMC sessions are logged off. Users with logged-off sessions must reconnect to the new active CMC.

About Standby CMC

The standby CMC is identical to and is maintained as a mirror of the active CMC. The active and standby CMCs must both be installed with the same firmware revision. If the firmware revisions differ, the system reports as "redundancy degraded".

The standby CMC assumes the same settings and properties of the active CMC. You must maintain the same firmware version on both the CMCs, but you do not have to duplicate configuration settings on the standby CMC.

NOTE: For information about installing a CMC, see the *VRTX Owner's Manual*. For instructions about installing CMC firmware on your standby CMC, see Updating Firmware.

CMC Failsafe Mode

The PowerEdge VRTX enclosure enables the fail-safe mode to protect the servers and I/O module from not functioning. The fail-safe mode is enabled when a CMC is not in control of the chassis. During the CMC failover period, or during a single CMC management loss:

- You cannot turn on the newly installed servers.
- You cannot remotely access existing servers.
- Server performance reduces to limit power consumption until management of the CMC is restored.

The following are some of the conditions that can result in CMC management loss:

- CMC removal Chassis management resumes after replacing CMC, or after failover to standby CMC.
- CMC network cable removal or network connection loss Chassis management resumes after the chassis fails over to the standby CMC. Network failover is only enabled in redundant CMC mode.
- CMC reset Chassis management resumes after CMC reboots or chassis fails over to the standby CMC.
- CMC failover command issued Chassis management resumes after the chassis fails over to the standby CMC.
- CMC firmware update Chassis management resumes after CMC reboots or chassis fails over to the standby CMC. It is recommended that you update the standby CMC first so that there is only one failover event.
- CMC error detection and correction Chassis management resumes after CMC resets or chassis fails over to the standby CMC.
- NOTE: You can configure the enclosure either with a single CMC or with redundant CMCs. In redundant CMC configurations, if the primary CMC loses communication with the enclosure or the management network, the standby CMC takes over the chassis management.

Active CMC Election Process

There is no difference between the two CMC slots; that is, slot does not indictate precedence. Instead, CMC that is installed or started first, assumes the role of an active CMC. If AC power is applied with two CMCs installed, CMC installed in CMC chassis slot 1 normally assumes the active role. The active CMC is indicated by a blue LED.

If two CMCs are inserted into a chassis that is already turned on, automatic active- or standby negotiation can take upto two minutes. Normal chassis operation resumes when the negotiation is complete.

Obtaining Health Status of Redundant CMC

You can view the health status of the standby CMC in the web interface. For more information about accessing CMC health status in the web interface, see Viewing Chassis Information and Monitoring Chassis and Component Health.

Configuring Front Panel

You can configure the following:

- Power button
- LCD
- DVD drive

Configuring Power Button

To configure the chassis power button:

- 1. In the left pane, click Chassis Overview > Front Panel > Setup.
- 2. On the Front Panel Configuration page, under the Power Button Configuration section, select the Disable Chassis Power Button option, and then click Apply.

The chassis power button is disabled.

Configuring LCD

- 1. In the left pane, click Chassis Overview > Front Panel > Setup .
- 2. On the Configuration page, under the LCD Configuration section:
 - Select the Lock Control Panel LCD option to disable any configuration that you can perform using the LCD interface.
 - From the LCD Language drop-down menu, select the required language.
 - From the LCD Orientation drop-down menu, select the required mode Tower Mode or Rack Mode.
 - NOTE: When you configure the chassis by using the LCD wizard, if you select the Auto-Apply settings to newly inserted servers option, you cannot disable the Auto-Apply settings to newly inserted servers function by using a basic license. If you do not want the function to take effect, either ignore the message displayed on the LCD, which will automatically disappear; or, press the Do not accept button on the LCD, and then push the center button.
- 3. Click Apply.

Accessing a Server Using KVM

To map the server to the KVM and enable accessing the server remote console through the KVM interface, you can use the CMC web interface, RACADM, or the LCD interface.

Mapping a Server to KVM Using CMC Web Interface

Make sure the KVM console is connected to the chassis.

To map a server to a KVM:

- 1. In the left pane, click Chassis Overview > Front Panel > Setup .
- 2. On the Front Panel Configuration page, under the KVM Configuration section, from the KVM Mapped list, select the slot that must be mapped to a KVM, and then click Apply.
- NOTE: The KVM allows mapping to all the server slots. Inserting a full-height or replacing a half-height server with full-height server does not change the mapping behavior. However, if the KVM is mapped to a lower slot and the slot has a full-height server, then the KVM is available only through the upper slot. You must remap the KVM to the upper slots.

Mapping the Server to KVM Using LCD

Make sure the KVM console is connected to the chassis.

To map the server to the KVM using the LCD — From the **Main Menu** screen on the LCD, go to **KVM Mapping**, select the server that must be mapped, and then press OK.

Mapping a Server to a DVD Drive

To map the server to the chassis DVD drive:

- 1. In the left pane, click Chassis Overview > Front Panel > Setup .
- On the Front Panel Configuration page, under the DVD Drive Configuration section:
 From the DVD Mapped drop-down menu, select one of the servers. Select the servers for which chassis DVD drive access is required.
- 3. Click Apply.

The DVD allows mapping to all the server slots. Inserting a full-height or replacing a half-height server with full-height server does not change the mapping behavior. However, if the DVD is mapped to a lower slot and the slot has a full-height server, then the DVD is available only through the upper slot. You must remap the DVD to the upper slots.

Logging in to CMC

You can log in to CMC as a CMC local user, as a Microsoft Active Directory user, or as an LDAP user. The default user name and password is root and calvin respectively. You can also log in using Single Sign-On or a Smart Card.

NOTE: CMC does not support the following special characters as user name or password from chassis profile using XML:

Topics:

- Accessing CMC Web Interface
- Logging in to CMC as a Local User, Active Directory User, or LDAP User
- Logging in to CMC Using a Smart Card
- Logging in to CMC Using Single Sign-on
- Logging In To CMC Using Serial, Telnet, Or SSH Console
- Accessing CMC Using RACADM
- Logging in to CMC Using Public Key Authentication
- Multiple CMC Sessions
- Changing Default Login Password
- Enabling or Disabling Default Password Warning Message
- Use case scenarios

Accessing CMC Web Interface

Before you log in to CMC using the web interface, make sure that you have configured a supported web browser (Internet Explorer or Firefox) and the user account is created with the required privileges.

NOTE: If you are using Microsoft Internet Explorer, connect using a proxy, and if you see the error The XML page cannot be displayed, you must disable the proxy to continue.

To access the CMC web interface:

- 1. Open a web browser supported on your system.
 - For the latest information on supported web browsers, see the *Dell Systems Software Support Matrix* located at **dell.com/support/manuals**.
- 2. In the Address field, type the following URL, and then press <Enter>:
 - To access CMC using IPv4 address: https://<CMC IP address>
 - If the default HTTPS port number (port 443) was changed, type: https://<CMC IP address>:<port number>
 - To access CMC using IPv6 address: https://[<CMC IP address>]
 - If the default HTTPS port number (port 443) was changed, type: https://[<CMC IP address>]:<port number>, where <CMC IP address> is the IP address for CMC and <port number> is the HTTPS port number.

The CMC Login page appears.

i) NOTE: While using IPv6, you must enclose the CMC IP address in parenthesis ([]).

Logging in to CMC as a Local User, Active Directory User, or LDAP User

To log in to CMC, you must have a CMC account with the **Log In to CMC** privilege. The default CMC user name is root, and the password is calvin. The root account is the default administrative account that ships with CMC.

(i) NOTE:

- For added security, it is strongly recommended that you change the default password of the root account during initial set up.
- When Certificate Validation is enabled, FQDN of the system should be provided. If certificate validation is enabled and IP address is provided for the Domain Controller, then the login is not successful.

CMC does not support extended ASCII characters, such as ß, å, é, ü, or other characters used primarily in non-English languages.

To log in as a local user, Active Directory user, or LDAP user.

- 1. In the **Username** field, type your user name:
 - CMC user name: <user name>
 - Active Directory user name: <domain>\<user name>, <domain>/<user name> or <user>@<domain>.
 - LDAP user name: <user name>
 - i NOTE: This field is case-sensitive.
- 2. In the Password field, type the user password.
 - i NOTE: For Active Directory user, the Username field is case-sensitive.
- 3. From the drop-down menu of the **Domain** field, select the required domain.
- **4.** Optionally, select a session timeout. This is the duration for which you can stay logged in with no activity before you are automatically logged out. The default value is the **Web Service Idle Timeout.**
- 5. Click OK.

You are logged into CMC with the required user privileges.

You cannot log in to the Web interface with different user names in multiple browser windows on a single workstation.

NOTE: If LDAP authentication is enabled and you attempt logging into CMC using the local credentials, the credentials are first checked in the LDAP server and then in CMC.

Logging in to CMC Using a Smart Card

To use this feature, you must have an Enterprise License. You can log in to CMC using a smart card. Smart cards provide Two Factor Authentication (TFA) that provide two-layers of security:

- Physical smart card device.
- Secret code such as a password or PIN.

Users must verify their credentials using the smart card and the PIN.

NOTE: You cannot use the IP address to log in to CMC using the Smart Card login. Kerberos validates your credentials based on the Fully Qualified Domain Name (FQDN).

Before you log in as an Active Directory user using a Smart Card, make sure to:

- Upload a Trusted Certificate Authority (CA) certificate (CA-signed Active Directory certificate) to CMC
- Configure the DNS server.
- Enable Active Directory login.
- Enable Smart Card login.

To log in to CMC as an Active Directory user using a smart card:

Log in to CMC using the link https://cmcname.domain-name>.
 The CMC Login page is displayed asking you to insert a smart card.

- NOTE: If you changed the default HTTPS port number (port 80), access the CMC web page using <cmcname.domain-name>:<port number>, where cmcname is the CMC host name for CMC, domain-name the domain name, and port number is the HTTPS port number.
- 2. Insert the smart card and click **Login**. The PIN dialog box is displayed.
- 3. Type the PIN and click Submit.
 - NOTE: If the smart card user is present in Active Directory, an Active Directory password is not required. Else, you have to log in by using an appropriate username and password.

You are logged in to CMC with your Active Directory credentials.

Logging in to CMC Using Single Sign-on

When Single Sign-On (SSO) is enabled, you can log in to CMC without providing your domain user authentication credentials, such as user name and password. To use this feature, you must have an Enterprise License.

NOTE: You cannot use the IP address to log in to the SSO. Kerberos validates your credentials against the Fully Qualified Domain Name (FQDN).

Before logging in to CMC using SSO, make sure that:

- You have logged in to your system using a valid Active Directory user account.
- Single Sign-On option is enabled during the Active Directory configuration.

To log in to CMC using SSO:

- 1. Log in to the client system using your network account.
- 2. Access the CMC web interface by using: https://<cmcname.domain-name>
 For example, cmc-6G2WXF1.cmcad.lab,, where cmc-6G2WXF1 is the cmc-name and cmcad.lab is the domain name.
 - NOTE: If you have changed the default HTTPS port number (port 80), access the CMC web interface using <cmcname.domain-name>:<port number>, where the *cmcname* is the CMC host name for CMC, **domain-name** is the domain name, and **port number** is the HTTPS port number.

CMC logs you in, using the Kerberos credentials that were cached by your browser when you logged in using your valid Active Directory account. If the login is unsuccessful, the browser is redirected to the normal CMC login page.

NOTE: If you are not logged in to the Active Directory domain and are using a browser other than Internet Explorer, the login is unsuccessful and the browser displays a only blank page.

Logging In To CMC Using Serial, Telnet, Or SSH Console

You can log in to CMC through a serial, Telnet, or SSH connection.

After you configure your management station terminal emulator software and managed node BIOS, perform the following tasks to log in to CMC:

- 1. Connect to CMC using your management station terminal emulation software.
- 2. Type your CMC user name and password, and then press <Enter>. You are logged in to CMC.

Related information

Using Telnet console with CMC on page 139 Configuring Linux Minicom on page 143

Accessing CMC Using RACADM

RACADM provides a set of commands that allow you to configure and manage CMC through a text-based interface. RACADM can be accessed using a Telnet/SSH or serial connection, using the Dell CMC console on the KVM, or remotely using the RACADM command line interface installed on a management station.

The RACADM interface is classified as:

- Remote RACADM Allows you to run RACADM commands on a management station with the -r option and the DNS name or IP address of the CMC.
 - NOTE: Remote RACADM is included on the *Dell Systems Management Tools and Documentation DVD* and is installed on a management station.
- Firmware RACADM Allows you to log in to the CMC using Telnet, SSH, or a serial connection. With firmware RACADM, you can the RACADM implementation that is part of the CMC firmware.

You can use remote RACADM commands in scripts to configure multiple CMCs. You cannot run the scripts directly on the CMC web interface, because CMC does not support it.

For more information about RACADM, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

For more information about configuring multiple CMCs, see Configuring Multiple CMCs Using RACADM.

Logging in to CMC Using Public Key Authentication

You can log in to the CMC over SSH without typing a password. You can also send a single RACADM command as a command line argument to the SSH application. The command line options behave similar to the remote RACADM, because the session ends after the command is completed.

Before logging in to CMC over SSH, make sure that the public keys are uploaded. To use this feature, you must have an Enterprise License.

For example:

- Logging in: ssh service@<domain> or ssh service@<*IP_address*>, where IP_address is the CMC IP address.
- Sending RACADM commands: ssh service@<domain> racadm getversion and ssh service@<domain> racadm getsel

When you log in using the service account, if a passphrase was set up when creating the public or private key pair, you may be prompted to enter that passphrase again. If the passphrase is used with the keys, client systems running Windows and Linux provide methods to automates the method. On client systems running Windows, you can use the Pageant application. It runs in the background and makes entering the passphrase transparent. For client systems running Linux, you can use the ssh agent. For setting up and using either of these applications, see their product documentation.

Multiple CMC Sessions

A list of multiple CMC sessions that are possible by using the various interfaces is given here.

Table 11. Multiple CMC sessions

Interface	Number of Sessions
CMC web interface	4
RACADM	4
Telnet	4
SSH	4

Changing Default Login Password

The warning message that prompts you to change the default password is displayed if:

- You log in to CMC with Configure Users privilege.
- Default password warning feature is enabled.
- Default user name and password for any currently enabled account are root and calvin respectively.

The same warning message is displayed if you log in using Active Directory or LDAP. Active Directory and LDAP accounts are not considered when determining if any (local) account has root and calvin as the credentials. A warning message is also displayed when you log in to CMC using SSH, Telnet, remote RACADM, or the Web interface. For Web interface, SSH, and Telnet, a single warning message is displayed for each session. For remote RACADM, the warning message is displayed for each command.

To change the credentials, you must have **Configure Users** privilege.

NOTE: A CMC log message is generated if the **Do not show this warning again** option is selected on the CMC **Login** page.

Changing Default Login Password Using Web Interface

When you log in to the CMC web interface, if the **Default Password Warning** page is displayed, you can change the password. To do this:

- 1. Select the Change Default Password option.
- 2. In the New Password field, type the new password.

The maximum characters for the password are 20. The characters are masked. The following characters are supported:

- 0-9
- A-Z
- a-z
- Special characters: +, &, ?, >, -, }, |, ., !, (, ', ,, _,[, ", @, #,), *, ;, \$,], /, §, %, =, <, :, {, !, \
- 3. In the Confirm Password field, type the password again.
- 4. Click Continue. The new password is configured and you are logged in to CMC.
 - NOTE: Continue is enabled only if the passwords provided in the New Password and Confirm Password fields match.

For information about the other fields, see the Online Help.

Changing Default Login Password Using RACADM

To change the password, run the following RACADM command:

```
racadm config -g cfgUserAdmin -o cfgUserAdminPassword -i <index> <newpassword>
```

where, <index> is a value from 1 to 16 (indicates the user account) and <newpassword> is the new user-defined password.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Enabling or Disabling Default Password Warning Message

You can enable or disable the display of the default password warning message. To do this, you must have **Configure Users** privilege.

Enabling or Disabling Default Password Warning Message Using Web Interface

To enable or disable the display of the default password warning message after logging in to iDRAC:

- Go to Chassis Controller > User Authentication > Local Users .
 The Users page is displayed.
- 2. In the **Default Password Warning** section, select **Enable**, and then click **Apply** to enable the display of the **Default Password Warning** page when you log in to CMC. Else, select **Disable**.

Alternatively, if this feature is enabled and you do not want to display the warning message for subsequent login operations, on the **Default Password Warning** page, select the **Do not show this warning again** option, and then click **Apply**.

Enabling or Disabling Warning Message to Change Default Login Password Using RACADM

To enable the display of the warning message to change the default login password using RACADM, use racadm config -g cfgRacTuning -o cfgRacTuneDefCredentialWarningEnable<0> or <1> object. For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Use case scenarios

This section describes typical use cases and tasks that you can perform with Chassis Management Controller Version 3.0 for Dell PowerEdge VRTX.

Conversion of External Shared PERC 8 card High Availability to Non-High Availability Mode using Web Interface

Dell PowerEdge VRTX chassis must have 2 External Shared PERC 8 cards in PCI slot 5 and PCI slot 6 in HA mode.

Workflow

- 1. Power down chassis. Disconnect all SAS cables from External Shared PERC 8 cards to MD12x0 enclosures.
- 2. Power Up chassis.
- 3. Login to CMC web interface and navigate to the Storage→ Controllers→ Troubleshooting and disable Fault tolerance from the drop-down menu for External Shared PERC 8 card in slot 5, click Apply and select disable for slot 6, and then click Apply.
- 4. Resetting both the PERCs may take two minutes to reflect in Non-HA mode.
- 5. Power down chassis and connect the enclosures in Non-HA Mode.
- 6. Power Up chassis.
- 7. External Shared PERC 8 card is not in HighAvailabiltiy Mode and navigate to **Storage→ Troubleshooting→ Setup**Troubleshooting to view the Non-HA status.

Conversion of External Shared PERC 8 card Non-High Availability to High Availability Mode using Web Interface

Dell PowerEdge VRTX chassis must have 2 External Shared PERC 8 cards in PCI slot 5 and PCI slot 6.

Workflow

- 1. Power down chassis. Disconnect all SAS cables from External Shared PERC 8 cards to MD12x0 enclosures.
- 2. Power Up chassis.
- 3. Login to CMC web interface and navigate to the Storage→ Controllers→ Troubleshooting and enable Fault tolerance from the drop-down menu for External Shared PERC 8 card in slot 5, click Apply and select disable for slot 6, and then click Apply.
- 4. Resetting both the PERCs may take two minutes to reflect in HA mode.
- 5. Power down chassis and connect the enclosures in HA Mode.
- 6. Power Up chassis.
- 7. External Shared PERC 8 card is in HighAvailabiltiy Mode and navigate to **Storage→ Troubleshooting→ Setup**Troubleshooting to view the HA status.

Conversion of External Shared PERC 8 card High Availability to Non-High Availability Mode using RACADM

Dell PowerEdge VRTX chassis must have 2 External Shared PERC 8 cards in PCI slot 5 and PCI slot 6 must be in HA mode.

Workflow

- 1. Power down chassis. Disconnect all SAS cables from External Shared PERC 8 cards to MD12x0 enclosures.
- 2. Power Up chassis.
- 3. Login to CMC Racadm and run the following command when the servers are in powered off state:

```
racadm raid set controllers:RAID.ChassisSlot.5-1 -p HighAvailabilityMode None
```

- 4. Run the command racadm raid set controllers:RAID.ChassisSlot.6-1 -p HighAvailabilityMode None on External Shared PERC 8 card in slot 6.
- 5. Resetting both the PERCs may take two minutes to reflect in HA mode.
- 6. Power down chassis and connect the enclosures in Non-HA Mode.
- 7. Power Up chassis.
- 8. External Shared PERC 8 card is not in HighAvailabiltiy Mode and the following command is used to view the status:

racadm raid get controllers -o -p HighAvailabilityMode

Conversion of External Shared PERC 8 card Non-High Availability to High Availability Mode using RACADM

Dell PowerEdge VRTX chassis must have External Shared PERC 8 cards in PCI slot 5 and PCI slot 6.

Workflow

- 1. Power down chassis. Disconnect all SAS cables from External Shared PERC 8 cards to MD12x0 enclosures.
- 2. Power Up chassis.
- 3. Login to CMC Racadm and run the following command when the servers are in powered off state:

```
racadm raid set controllers:RAID.ChassisSlot.5-1 -p HighAvailabilityMode ha
```

- 4. Run the command racadm raid set controllers:RAID.ChassisSlot.6-1 -p HighAvailabilityMode ha on External Shared PERC 8 card in slot 6.
- 5. Resetting both the PERCs may take two minutes to reflect in HA mode.
- 6. Power down chassis and connect the enclosures in HA Mode.
- 7. Power Up chassis.
- 8. External Shared PERC 8 card is in HighAvailabiltiy Mode and the HA status is seen by using the following command:

racadm raid get controllers -o -p HighAvailabilityMode

Updating Firmware

You can update firmware for:

- CMC
- Chassis infrastructure
- VRTX Expander or Storage Backplane Expander Firmware of Inegrated or external enclosures
- Physical Disks (HDD) per enclosure
 - i NOTE: You can update the HDD firmware only if required.

You can update firmware for the following I/O and server components:

- I/O Module
- BIOS
- iDRAC
- Lifecycle Controller
- 32-bit diagnostics
- Operating System Drivers Pack
- Network Interface Controllers
- RAID controllers on the server module

i NOTE: Firmware update may take several minutes to complete.

Topics:

- Downloading CMC Firmware
- Viewing Currently Installed Firmware Versions
- Updating the CMC Firmware
- Updating Chassis Infrastructure Firmware
- Updating Server iDRAC Firmware
- Updating Server Component Firmware
- Viewing Firmware Inventory
- Saving Chassis Inventory Report Using CMC Web Interface
- Configuring Network Share Using CMC Web Interface
- Lifecycle Controller Job Operations
- Rolling Back Server Component Firmware
- Upgrading Server Component Firmware
- Deleting Scheduled Server Component Firmware Jobs
- Updating Storage Component Using CMC Web Interface
- Recovering iDRAC Firmware Using CMC

Downloading CMC Firmware

Before beginning the firmware update, download the latest firmware version from **support.dell.com**, and save it to your local system.

While updating VRTX Chassis firmware, it is recommended to update the firmware versions of the chassis components in the following order:

- 1. Blade components firmware
- 2. CMC firmware
- 3. Chassis infrastructure firmware
- **4.** Shared PERC8 firmware (integrated and external)
- 5. Internal Storage Backplane firmware and external enclosure's expanders
- **6.** HDD firmware (external and integrated enclosures)

For more information about the update sequence for VRTX chassis, see the CMC Firmware 3.0 Release Notes available at dell.com/cmcmanuals.

Viewing Currently Installed Firmware Versions

You can view the currently installed firmware versions using the CMC web interface or RACADM.

Viewing Currently Installed Firmware Versions Using CMC Web Interface

In the CMC web interface, go to any of the following pages to view the current firmware versions:

- Chassis Overview > Update
- Chassis Overview > Chassis Controller > Update
- Chassis Overview > Server Overview > Server Component Update
- Chassis Overview > I/O Module Overview > Update
- Chassis Overview > Storage > Storage Component Update

The **Firmware Update** page displays the current version of the firmware for each listed component and allows you to update the firmware to the latest version.

If the chassis contains an earlier generation server, whose iDRAC is in recovery mode or if CMC detects that iDRAC has corrupted firmware, then the earlier generation iDRAC is also listed on the **Firmware Update** page.

Viewing Currently Installed Firmware Versions Using RACADM

To view the IP information for iDRAC and CMC, and the CMC service- or asset tag using RACADM, run the racadm getsysinfo sub-command. For more information about other RACADM commands, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Updating the CMC Firmware

You can update the CMC firmware using web interface or RACADM. The firmware update, by default, retains the current CMC settings. During the update process, you can reset CMC configuration settings back to the factory-default settings.

i NOTE: To update firmware on CMC, you must have the Chassis Configuration Administrator privilege.

If a Web user interface session is used to update system component firmware, the **Idle Timeout (0, 60–10800)** setting must be set to a higher value to accommodate the file transfer time. In some cases, the firmware file transfer time may be as high as 30 minutes. To set the idle timeout value, see Configuring Services.

During CMC firmware updates, it is normal for some or all of the fan units in the chassis to rotate at 100% speed.

If you have redundant CMCs installed in the chassis, it is recommended to update both the CMCs to the same firmware version, at the same time, with a single operation. If CMCs have different firmware and a failover occurs, unexpected results may occur.

(i) NOTE:

- The CMC firmware cannot be updated to any earlier version other than 2.0 for a chassis that is configured with 1600W PSII
- CMC firmware update or roll back is supported only for firmware versions 1.2, 1.25, 1.3, 1.31, 1.35, 1.36, 2.0, 2.01 and 2.04 later. For any version other than these, first update to any of these versions, and then update to the required version.

The Active CMC resets and becomes temporarily unavailable after the firmware has been uploaded successfully. If a standby CMC is present, the standby and active roles swap. The standby CMC becomes the active CMC. If an update is applied only to the active CMC, after the reset is complete, the active CMC does not run the updated image, only the standby CMC has that image. In general, it is highly recommended to maintain identical firmware versions for the active and standby CMCs.

When the standby CMC has been updated, swap the CMCs' roles so that the newly updated CMC becomes the active CMC and the CMC with the earlier firmware becomes the standby. For information about swapping roles, see the cmcchangeover

command section in the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide. Running this command helps you to verify that the update has succeeded and that the new firmware is working properly, before you update the firmware in the second CMC. When both CMCs are updated, you can use the cmcchangeover command to restore the CMCs to their previous roles. CMC firmware revision 2.x updates both the primary CMC and the redundant CMC without running the cmcchangeover command.

During the final phases of the firmware update process in CMC, the browser session and connection with CMC is lost temporarily as the CMC is not connected to the network. CMC generates the chassis overall health as critical owing to the temporary network loss. When CMC restarts after a few minutes, log in to it. The CMC then generates the chassis overall health as healthy and the CMC network link is up. After CMC resets, the new firmware version is displayed on the **Firmware Update** page.

To avoid disconnecting other users during a reset, notify authorized users who may log in to CMC and check for active sessions on the **Sessions** page. To open the **Sessions** page, click **Chassis Overview** in the left pane, click **Network**, and then click the **Sessions**.

When transferring files to and from CMC, the file transfer icon spins during the transfer. If your icon is not animated, make sure that your browser is configured to allow animations. For more information about allowing animations in the browser, see Allow Animations in Internet Explorer.

Signed CMC Firmware Image

For VRTX CMC 2.0 and later, the firmware includes a signature. The CMC firmware verifies the signature to ensure the authenticity of the uploaded firmware. The firmware update process is successful only if the firmware image is authenticated by CMC to be a valid image from the service provider and has not been altered. The firmware update process is stopped if CMC cannot verify the signature of the uploaded firmware image. A warning event is then logged and an appropriate error message is displayed.

Signature verification can be performed on VRTX firmware versions 1.2 and later. For firmware downgrade to VRTX versions earlier than 1.2, first update the firmware to a VRTX CMC version that is later than or equal to 1.2, but earlier than 2.0. After this update, firmware downgrade to earlier, unsigned VRTX versions can be performed.

Updating CMC and Mainboard Firmware

The shared capabilities of External Shared PERC 8 card are not available until both CMC and mainboard firmware are updated.

(i) NOTE:

- To view the MD12x0 cabling diagram, refer to Upgrading PowerEdge VRTX to Support Shared Storage Expansion User's Guide or Dell Shared PowerEdge RAID Controller (PERC) 8 Cards For Dell PowerEdge VRTX Systems User's Guide available at dell.com/support/manuals.
- The external shared storage adapter requires that you update the CMC v2.20 or later and mainboard v2.21 or later to support External Shared PERC 8 card.
- You cannot downgrade the CMC firmware prior to 2.2 with external shared adapters.

To update the CMC and mainboard firmware:

- 1. Update CMC firmware.
- 2. Update mainboard firmware.
- 3. Power down the chassis and install shared storage adapters in PCle slot 5 and slot 6.
- 4. Power on chassis.
- 5. After powering on the chassis, update external shared storage adapters.
- NOTE: By default, the external shared External Shared PERC 8 card is in non-fault tolerant mode. It has to be changed to Fault Tolerant mode after it is cabled appropriately. For more information, refer *Upgrading PowerEdge VRTX to Support Shared Storage Expansion*.

In an event, if you want to rollback CMC or MPC/mainboard firmware or CMC and MPC firmware version, do the following tasks:

To Rollback the CMC and mainboard firmware:

- 1. Power off the chassis.
- 2. Remove all external storage adapters from the PCI slots.
- 3. Power on the chassis.

4. Roll back the CMC and/or mainboard firmware.

You cannot downgrade the CMC, if external shared storage adapter is detected.

If the processes are not followed in order, system behavior becomes random and parts of the system may become unstable. The CMC logs IOV or RAID controller messages. Only shared storage VA mappings for PERC 1 and PERC 2 are visible in the older version of CMC. All external shared storage VA mappings do not exist in the previous version of the CMC. If an External Shared PERC 8 card is inserted after the rollback, the CMC treats it as a non-shared adapter. It may happen that the HOST PERC driver does not support the External Shared PERC 8 card.

Updating CMC Firmware Using Web Interface

(i) NOTE:

- Before applying the CMC update, ensure that the chassis is powered on. If the blades are powered on, it is not necessary to power off the blades to perform the CMC update.
- Downgrading the CMC Firmware prior to 2.1 with external shared adapters is blocked.

To update the CMC firmware using the CMC web interface:

- 1. In the left pane, go to any of the following pages:
 - Chassis Overview > Update
 - Chassis Overview > Chassis Controller > Update
- 2. On the **Firmware Update** page, in the **CMC Firmware** section, select the required components under the **Update Targets** column for the CMC or CMCs (if standby CMC is present) you want to update, and then click **Apply CMC Update**.
- 3. In the **Firmware Image** field, click **Browse** (Internet Explorer or Firefox) or **Choose File** (Google Chrome) to browse through to the file location. The default name of the CMC firmware image file is vrtx_cmc.bin.
- **4.** Click **Begin Firmware Update**. The **Firmware Update Progress** section provides firmware update status information. A status indicator displays on the page while the image file is uploaded. File transfer time varies based on the connection speed. When the internal update process begins, the page automatically refreshes and the Firmware update timer is displayed.
- 5. For a standby CMC, when the update is complete, the **Update State** field displays **Done**. For an active CMC, during the final phases of the firmware update process, the browser session and connection with CMC is lost temporarily because the active CMC is not connected to the network. You must log in after a few minutes, when the active CMC has restarted. After CMC resets, the new firmware is displayed on the **Firmware Update** page.
 - NOTE: After the firmware update, delete the files from the web browser cache. For instructions about clearing the browser cache, see the web browser's online help.

Additional instructions:

- During a file transfer, do not click the **Refresh** icon or navigate to another page.
- To cancel the process, select the Cancel File Transfer and Update option. This option is available only during file transfer.
- The **Update State** field displays the firmware update status.
 - i NOTE: The update process may take several minutes for CMC.

Updating CMC firmware using RACADM

To update CMC firmware using RACADM, use the fwupdate subcommand. For more information about RACADM commands, see Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

(i) **NOTE:** Run the firmware update command through only one remote racadm session at a time.

Updating Chassis Infrastructure Firmware

The chassis infrastructure update operation updates components such as the Main Board and PCIe subsystem management firmware.

NOTE: To update the chassis infrastructure firmware, make sure that the chassis is turned on and the servers are turned off.

i) NOTE: When the mainboard is upgraded to a later version, the chassis and Chassis Management Controller may reboot.

Updating Chassis Infrastructure Firmware Using CMC Web Interface

- 1. Go to any of the following pages:
 - Chassis Overview > Update.
 - Chassis Overview > Chassis Controller > Update.
- 2. On the **Firmware Update** page, in the **Chassis Infrastructure Firmware** section, in the **Update Targets** column, select the option, and then click **Apply Chassis Infrastructure Firmware**.
- 3. On the Firmware Update page, click Browse, and then select the appropriate chassis infrastructure firmware.
- 4. Click Begin Firmware Update, and then click Yes.

The **Firmware Update Progress** section provides firmware update status information. While the image file uploads, a status indicator displays on the page. File transfer time varies on the basis of connection speed. When the internal update process begins, the page automatically refreshes and the firmware update timer is displayed.

Additional instructions to follow:

- Do not click the **Refresh** icon, or navigate to another page during the file transfer.
- The **Update State** field displays the firmware update status.

When the update is complete, there is a brief loss of connectivity to the main board, because it resets and the new firmware is displayed on the **Firmware Update** page.

Updating Chassis Infrastructure Firmware Using RACADM

To update chassis infrastructure firmware using RACADM, use the fwupdate sub-command. For more information about using the RACADM commands, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Updating Server iDRAC Firmware

You can update firmware for iDRAC using the CMC Web interface or the RACADM. To use this feature, you must have an Enterprise License.

The iDRAC firmware version must be 1.40.40 or later for servers with iDRAC.

The iDRAC (on a server) resets and is temporarily unavailable after a firmware update.

NOTE: To update an iDRAC firmware using the Chassis Management Controller, an SD card must be available in the chassis. However, to update the iDRAC firmware through the iDRAC Web interface, an SD card is not required in CMC. For more information on launching iDRAC Web interface from CMC, see Launching iDRAC from Server Status Page.

Updating Server iDRAC Firmware Using Web Interface

To update the iDRAC firmware in the server:

- 1. Go to any of the following pages:
 - Chassis Overview > Update.
 - Server Overview > Update > Server Component Update.

The Firmware Update page is displayed.

(i) NOTE:

You can also update server iDRAC firmware using **Chassis Overview** > **Server Overview** > **Update**. For more information, see Updating Server Component Firmware.

2. To update the iDRAC7 or iDRAC8 firmware, in the iDRAC7 Firmware or iDRAC8 Firmware section respectively, click the Update link of the server for which you want to update the firmware.

The Server Component Update page is displayed. To continue, see Updating Server Component Firmware.

- 3. In the **Firmware Image** field, enter the path to the firmware image file on the management station or shared network, or click **Browse** to navigate to the file location. The default iDRAC firmware image name is firming.imc.
- 4. Click Begin Firmware Update, and then click Yes.

The **Firmware Update Progress** section provides firmware update status information. A progress bar indicates the status of the upload process. File transfer time varies on the basis of connection speed. When the internal update process begins, the page automatically refreshes and the firmware update timer is displayed.

- NOTE: Additional instructions to follow:
 - Do not click the **Refresh** icon or navigate to another page during the file transfer.
 - To cancel the process, click Cancel File Transfer and Update. This option is available only during file transfer.
 - The **Update State** field displays the firmware update status.

It may take up to 10 minutes to update the iDRAC firmware.

Updating Server Component Firmware

The one-to-many update feature in CMC enables you to update server component firmware across multiple servers. You can update the server components using the Dell Update Packages available on the local system or on a network share. This operation is enabled by leveraging the Lifecycle Controller functionality on the server.

The Lifecycle Controller service is available on each server and is facilitated by iDRAC. You can manage the firmware of the components and devices on the servers using the Lifecycle Controller service. The Lifecycle Controller uses an optimization algorithm to update the firmware that efficiently reduces the number of restarts.

The Lifecycle Controller provides module update support for iDRAC7 and later servers.

- NOTE: Before using the Lifecycle Controller-based update feature, server firmware versions must be updated. You must also update the CMC firmware before updating the server component firmware modules.
- NOTE: To update component firmware, the CSIOR option must be enabled for servers. To enable CSIOR on:
 - 12th generation servers and later— After restarting the server, from the F2 setup, select iDRAC Settings > Lifecycle Controller, enable CSIOR and save the changes.
 - 13th generation servers After rebooting the server, when prompted, press F10 to access Lifecycle Controller. Go
 to the Hardware Inventory page by selecting Hardware Configuration > Hardware Inventory. On the Hardware
 Inventory page, click Collect System Inventory on Restart.

The **Update from File** method enables you to update the server component firmware using DUP files stored on a local system. You can select the individual server components to update the firmware using the required DUP files. You can update large number of components at a time by using an SD Card to store DUP file of more than 48 MB memory size.

- NOTE: Note the following:
 - While selecting the individual server components for update, make sure that there are no dependencies between the selected components. Else, selecting some components that have dependencies on other components for update may cause the server to stop functioning abruptly.
 - Make sure to update the server components in the recommended order. Else, the process of component firmware update may become unsuccessful.

Always update the server component firmware modules in the following order:

- o BIOS
- o Lifecycle Controller
- o iDRAC

The Single Click all blade update or the **Update from Network Share** method enables you to update the server component firmware using DUP files stored on a network share. You can use the Dell Repository Manager (DRM) based update feature to access the DUP files stored on a network share and update the server components in a single operation. You can set up a custom remote repository of firmware DUPs and binary images using the Dell Repository Manager and share it on the Network Share. Alternatively, use the Dell Repository Manager (DRM) to check for the latest available firmware updates. The Dell Repository Manager (DRM) ensures that the Dell systems are up-to-date with the latest BIOS, driver, firmware, and software. You can search for the latest updates available from the Support site (**support.dell.com**) for supported platforms based on Brand and Model or a Service Tag. You can download the updates or build a repository from the search results. For more

information on using the DRM to search for latest firmware updates, see Using Dell Repository Manager to Search for the Latest Updates on the Dell Support Site on the Dell Tech Center. For information on saving the inventory file that DRM uses as input to create the repositories, see Saving Chassis Inventory Report Using CMC Web Interface

- NOTE: The Single Click all blade update method has the following benefits:
 - Enables you to update all the components on all the blade servers with minimal clicks.
 - All the updates are packaged in a directory. This avoids individual upload of each component's firmware.
 - Faster and consistent method of updating the server components
 - Enables you to maintain a standard image with the required updates versions of the server components that can be used to update multiple servers in a single operation.
 - You can copy the directories of updates from the Dell Server Update Utility (SUU) download DVD or create and customize the required update versions in the Dell Repository Manager (DRM). You do not need the latest version of the Dell Repository Manager to create this directory. However, Dell Repository Manager version 1.8 provides an option to create a repository (directory of updates) based on the inventory that was exported from the servers in the chassis. For information on creating a repository using the Dell Repository Manager see the Dell Repository Manager Data Center Version 1.8 User's Guide and the Dell Repository Manager Business Client Version 1.8 User's Guide available at dell.com/support/manuals.

Lifecycle Controller provides module update support through iDRAC. It is recommended to update the CMC firmware before updating the server component firmware modules. After updating the CMC firmware, in the CMC Web interface, you can update the server component firmware on the **Chassis Overview** > **Server Overview** > **Update** > **Server Component Update** page. It is also recommended to select all the component modules of a server to be updated together. This enables Lifecycle Controller to use its optimized algorithms to update the firmware, reducing the number of reboots.

To update the server component firmware, using the CMC Web interface, click **Chassis Overview > Server Overview > Update > Server Component Update**.

If the server does not support the Lifecycle Controller service, the **Component/Device Firmware Inventory** section displays **Not Supported**. For the latest generation servers, install the Lifecycle Controller firmware and update the iDRAC firmware to enable the Lifecycle Controller service on the server. For earlier generation servers, this upgrade is not possible.

Normally, the Lifecycle Controller firmware is installed using an appropriate installation package that is executed on the server operating system. For supported servers, a special repair or installation package with an <code>.usc</code> file extension is available. This file enables you to install the Lifecycle Controller firmware through the firmware update facility available on the native iDRAC Web browser interface.

You can also install Lifecycle Controller firmware through an appropriate installation package executed on the server OS. For more information, see the *Dell Lifecycle Controller User's Guide*.

If Lifecycle Controller service is disabled on the server, the Component/Device Firmware Inventory section displays

Lifecycle Controller may not be enabled.

Server Component Update Sequence

In case of individual component updates, you must update the firmware versions for the server components in the following sequence:

- iDRAC
- Lifecycle Controller
- Diagnostics (optional)
- OS Driver Packs (optional)
- BIOS
- NIC
- RAID
- Other components
- NOTE: When you update the firmware versions for all the server components at one time, the update sequence is handled by Lifecycle Controller.

Enabling Lifecycle Controller

You can enable the Lifecycle Controller service when turning on a server:

- For iDRAC servers, on the boot console, to access **System Setup**, press the <F2> key.
- On the System Setup Main Menu page, go to iDRAC Settings > Lifecycle Controller, click Enabled. Go to the System Setup Main Menu page and click Finish to save the settings.

Canceling System Services enables you to cancel all scheduled jobs that are pending and remove them from the queue.

For more information about the Lifecycle Controller and supported server components, and device firmware management, see:

- Lifecycle Controller-Remote Services Quick Start Guide.
- delltechcenter.com/page/Lifecycle+Controller.

The **Server Component Update** page enables you to update various firmware components on the server. To use the features and functions on this page, you must have:

- For CMC: The Server Administrator privilege.
- For iDRAC: The Configure iDRAC privilege and Log in to iDRAC privilege.

In case of insufficient privileges, you can only view the firmware inventory of components and devices on the server. You cannot select any components or devices for any type of Lifecycle Controller operation on the server.

Choosing Server Component Firmware Update Type Using CMC Web Interface

To select the type of server component update type:

- In the system tree, go to Server Overview, and then click Update > Server Component Update.
 The Server Component Update page is displayed.
- 2. In the Choose Update Type section, select the required update method:
 - Update from File
 - Update from Network Share

Filtering Components for Firmware Updates

Information about all the components and devices across all servers is retrieved at one time. To manage this large amount of information, Lifecycle Controller provides various filtering mechanisms.

(i) NOTE: To use this feature, you must have an Enterprise License.

The Component/Device Update Filter section in the Server Component Update page that allows you to filter the information based on the component, is available only for the Update by File mode.

These filters enable you to:

- Select one or more categories of components or devices for easy viewing.
- Compare firmware versions of components and devices across the server.
- To narrow the category of a particular component or device based on types or models, automatically filter the selected components and devices.
 - NOTE: Automatic filtering feature is important while using the Dell Update Package (DUP). The update programming of a DUP can be based on the type or model of a component or device. The automatic filtering behavior is designed to minimize the subsequent selection decisions after an initial selection is made.

Following are some examples where the filtering mechanisms are applied:

- If the BIOS filter is selected, only the BIOS inventory of all the servers is displayed. If the set of servers consists of a number
 of server models, and a server is selected for BIOS update, the automatic filtering logic automatically removes all the other
 servers that do not match with the model of the selected server. This makes sure that the selection of the BIOS firmware
 update image (DUP) is compatible with the correct model of the server.
 - Sometimes, a BIOS firmware update image may be compatible across a number of server models. Such optimizations are ignored in case this compatibility is no longer true in the future.
- Automatic filtering is important for firmware updates of Network Interface Controllers (NIC) and RAID Controllers. These
 device categories have different types and models. Similarly, the firmware update images (DUP) may be available in
 optimized forms, where a single DUP may be programmed to update multiple types or models of devices of a given category.

Filtering Components for Firmware Updates Using CMC Web Interface

To filter the devices:

- 1. In the left pane, go to Server Overview, and then click Update .
- 2. On the Server Component Update page, in the Component/Device Update Filter section, select one or more of the following:
 - BIOS
 - iDRAC
 - Lifecycle Controller
 - 32-Bit Diagnostics
 - OS Driver Pack
 - Network I/F Controller
 - RAID Controller

The Component/Device Update Filter section is displayed only for the Update by File mode of firmware update.

The **Firmware Inventory** section displays only the associated components or devices across all servers present in the chassis. After you select an item from the drop-down menu, only the components or devices associated with the ones in the are list displayed.

After the filtered set of components and devices is displayed in the inventory section, further filtering may occur when a component or device is selected for update. For example, if the BIOS filter is selected, then the inventory section displays all the servers with only their BIOS component. If a BIOS component on one of the servers is selected, the inventory is further filtered to display the servers that match the model name of the selected server.

If a filter is not selected and a selection for update of a component or device is made on the inventory section, then the filter associated with that selection is automatically enabled. Further filtering may occur where the inventory section displays all the servers that have a match for the selected component in terms of model, type, or some form of identity. For example, if a BIOS component on one of the servers is selected for update, the filter is automatically set to the BIOS and the inventory section displays the servers that match the model name of the selected server.

Filtering Components for Firmware Updates Using RACADM

To filter components for Firmware Updates using RACADM, run the **getversion** command:

racadm getversion -l [-m <module>] [-f <filter>]

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Viewing Firmware Inventory

You can view the summary of the firmware versions for all components and devices for all servers currently present in the chassis along with their status.

NOTE: To use this feature, you must have an Enterprise License.

Viewing Firmware Inventory Using CMC Web Interface

To view the firmware inventory:

- 1. In the left pane, click Server Overview, and then click Update.
- 2. On the Server Component Update page, view the firmware inventory details in the Component/Device Firmware Inventory section. On this page, you can view the following information:
 - Servers that currently do not support the Lifecycle Controller service are listed as Not Supported. A hyperlink is
 provided to an alternative page, where you can directly update only the iDRAC firmware. This page supports only iDRAC
 firmware update and not any other component and device on the server. iDRAC firmware update is not dependent on the
 Lifecycle Controller service.
 - If the server is listed as **Not Ready**, it indicates that when the firmware inventory was retrieved, the iDRAC on the server was still initializing. Wait for the iDRAC to be fully operational, and then refresh the page to retrieve the firmware inventory again.

- If the inventory of components and devices does not reflect what is physically installed on the server, invoke the Lifecycle Controller when the server is in the boot process. This action helps to refresh the integrated components and devices information and allows you to verify the currently installed components and devices. The inventory does not reflect the component and device information accurately when:
 - The server iDRAC firmware is updated to newly introduce the Lifecycle Controller functionality to the server management.
 - o The new devices are inserted into the server.

To automate this action the iDRAC Settings Utility provides an option that can be accessed through the boot console:

- a. For iDRAC servers, on the boot console, to access **System Setup**, press <F2>.
- b. On the System Setup Main Menu page, click iDRAC Settings > Collect System Inventory on Restart, select Enabled, go back to the System Setup Main Menu page, and then click Finish to save the settings.
- Options to perform the various Lifecycle Controller operations such as Update, Rollback, Reinstall, and Job Deletion are
 available. Only one type of operation can be performed at a time. Components and devices that are not supported may
 be listed as part of the inventory, but do not permit Lifecycle Controller operations.

The following table displays the component and devices information on the server:

Table 12. Component and Devices Information

Field	Description
Slot	Displays the slot occupied by the server in the chassis. Slot numbers are sequential IDs, from 1 to 4 (for the four available slots in the chassis), that help to identify the location of the server in the chassis. When there are less than four servers occupying slots, only those slots populated by servers are displayed.
Name	Displays the name of the server in each slot.
Model	Displays the model of the server.
Component/ Device	Displays a description of the component or device on the server. If the column width is too narrow, the mouse-over tool provides a view of the description.
Current Version	Displays the current version of component or device on the server.
Rollback Version	Displays the rollback version of component or device on the server.
Job Status	Displays the job status of any operations that are scheduled on the server. The job status is continuously updated dynamically. If a job completion with state completed is detected, then the firmware versions for the components and devices on that server are automatically refreshed in case there has been a change of firmware version on any of the components or devices. An information icon is also presented adjacent to the current state, which provides additional information about the current job status. This information can be viewed by clicking or pausing the mouse over the icon.
Update	Click to select the component or device for firmware update on the server.

Viewing Firmware Inventory Using RACADM

To view firmware inventory using RACADM, use the getversion command:

racadm getversion -l [-m <module>] [-f <filter>]

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Saving Chassis Inventory Report Using CMC Web Interface

To save the chassis inventory report:

In the system tree, go to Server Overview, and then click Update > Server Component Update.
The Server Component Update page is displayed.

2. Click Save Inventory Report.

The Inventory.xml file is saved on an external system.

NOTE: The Dell Repository Manager Application uses the *Inventory.xml* file as an input to create a repository of updates for all the blades available in the chassis. This repository can be later exported to a network share. **Update from**Network Share mode of firmware update uses this network share to update the components of all the servers. You must have CSIOR enabled on the individual servers and save the chassis inventory report every time there is a change to the chassis hardware and software configuration.

Configuring Network Share Using CMC Web Interface

To configure or edit the Network Share location or credentials:

- In the CMC Web interface, in the system tree, go to Server Overview and then click Network Share.
 The Edit Network Share page is displayed.
 - NOTE: When you have same folder for chassis, server, and boot identity profile, you may see performance issues if there is more than 100 profiles.
- 2. In the Network Share Settings section, configure the following settings as required:
 - Protocol
 - IP Address or Host Name
 - Share Name
 - Update folder
 - File Name (optional)
 - NOTE: Entering the **File Name** is optional only when the default catalog file name is *catalog.xml*. If the catalog file name is changed then the new name must be entered in this field.
 - Profile Folder
 - Domain Name
 - User Name
 - Password

For more information, see the CMC Online Help.

- **3.** Click **Test Directory** to verify whether the directories are readable and writeable.
- 4. Click **Test Network Connection** to verify if the network share location is accessible.
- 5. Click **Apply** to apply the changes to the network share properties.
 - (i) NOTE:

Click Back to return to the Server Component Update page.

Lifecycle Controller Job Operations

(i) NOTE: To use this feature, you must have an Enterprise License.

You can perform Lifecycle Controller operations such as:

- Re-install
- Rollback
- Update
- Delete Jobs

Only one type of operation can be performed at a time. Components and devices that are not supported may be listed as part of the inventory, but do not permit Lifecycle Controller operations.

To perform the Lifecycle Controller operations, you must have:

- For CMC: Server Administrator privilege.
- For iDRAC: Configure iDRAC privilege and Log in to iDRAC privilege.

A Lifecycle Controller operation scheduled on a server may take 10 to 15 minutes to complete. The process involves several server reboots during which the firmware installation is performed, which also includes a firmware verification stage. You can view the progress of this process using the server console. If there are several components or devices that need to be updated on a server, you can consolidate all the updates into one scheduled operation thus minimizing the number of reboots required.

Sometimes, when an operation is in the process of being submitted for scheduling through another session or context, another operation is attempted. In this case, a confirmation message is displayed indicating the situation and the operation must not be submitted. Wait for the operation in process to complete and then submit the operation again.

Do not navigate away from the page after an operation is submitted for scheduling. If an attempt is made, a confirmation message is displayed allowing the intended navigation to be cancelled. Otherwise, the operation is interrupted. An interruption, especially during an update operation may cause the firmware image file upload to be terminated before proper completion. After an operation has been submitted for scheduling, ensure that the confirmation message indicating that the operation has been successfully scheduled is acknowledged.

Reinstalling Server Component Firmware

You can reinstall the firmware image of the currently installed firmware for selected components or devices across one or more servers. The firmware image is available within the Lifecycle Controller.

Re-installing Server Component Firmware Using Web Interface

To reinstall a server component firmware:

- 1. In the left pane, click Server Overview > Update.
- 2. On the Server Component Update page, in the Choose Update Type section, select Update from File.
- 3. In the Current Version column, select the option for the component or device for which you want to reinstall the firmware.
- **4.** Select one of the following options:
 - Reboot Now Restart the server immediately.
 - On Next Reboot Manually restart the server at a later time.
- 5. Click Reinstall. The firmware version is reinstalled for the selected component or device.

Rolling Back Server Component Firmware

You can install the firmware image of the previously installed firmware for selected components or devices across one or more servers. The firmware image is available within the Lifecycle Controller for a rollback operation. The availability is subject to the version compatibility logic of the Lifecycle Controller. It also assumes the previous update was facilitated by the Lifecycle Controller.

(i) NOTE: To use this feature, you must have an Enterprise License.

Rolling Back Server Component Firmware Using the CMC Web Interface

To roll back the server component firmware version to an earlier version:

- 1. In the left pane, click Server Overview→ Update.
- 2. On the Server Component Update page, in the Choose Update Type section, select Update from File.
- 3. In the **Rollback Version** column, select the option for the component or device for which you want to roll back the firmware.
- 4. Select one of the following options:
 - Reboot Now Restart the server immediately.
 - On Next Reboot Manually restart the server at a later time.
- 5. Click Rollback. The previously installed firmware version is reinstalled for the selected component or device.

Upgrading Server Component Firmware

You can install the next version of the firmware image for selected components or devices across one or more servers. The firmware image is available within the Lifecycle Controller for a rollback operation. To use this feature, you must have an Enterprise License.

NOTE: For iDRAC and Operating System Driver packs firmware update, make sure the **Extended Storage** feature is enabled.

It is recommended to clear the job queue before initializing a server component firmware update. A list of all jobs on the server is available on the **Lifecycle Controller Jobs** page. This page enables deletion of single or multiple jobs or purging of all jobs on the server

BIOS updates are specific to the model of the server. Sometimes, even though a single Network Interface Controller (NIC) device is selected for firmware update on a server, the update may get applied to all the NIC devices on the server. This behavior is inherent in the Lifecycle Controller functionality and particularly the programming contained with the Dell Update Package (DUP). Currently, Dell Update Packages (DUP) that are less than 48 MB in size are supported.

If the update file image size is greater, the job status indicates that the download has failed. If multiple server component updates are attempted on a server, the combined size of all the firmware update files may also exceed 48 MB. In such a case, one of the component updates fails as its update file is truncated. To update multiple components on a server, it is recommended to update the Lifecycle Controller and 32-Bit Diagnostics components together first. These do not require a server reboot and are relatively quick to complete. The other components can then be updated together.

All Lifecycle Controller updates are scheduled for immediate execution. However, the system services can delay this execution sometimes. In such situations, the update fails as a result of the remote share that is hosted by the CMC being no longer available.

Upgrading Server Component Firmware From File Using CMC Web Interface

To upgrade the server components firmware version to the next version using the **Update from File** method:

- 1. In the CMC Web interface, in the system tree, go to **Server Overview** and then click **Update** > **Server Component Update**.
 - The Server Component Update page is displayed.
- 2. In the **Choose Update Type** section, select **Update from File**. For more information, see Choosing Server Component Firmware Update Type Using CMC Web Interface
- **3.** In the **Component/Device Update Filter** section, filter the component or device (optional). For more information see Filtering Components for Firmware Updates Using CMC Web.
- 4. In the **Update** column, select the checkbox(es) for the component or device for which you want to update the firmware to the next version. Use the CRTL key shortcut to select a type of component or device for update across all the applicable servers. Pressing and holding the CRTL key highlights all the components in yellow. While the CRTL key is pressed down, select the required component or device by enabling the associated check box in the **Update** column.

A second table is displayed that lists the selected type of component or device and a selector for the firmware image file. For each type of component, one selector for the firmware image file is displayed.

Few devices such as Network Interface Controllers (NICs) and RAID Controllers contain many types and models. The update selection logic automatically filters the relevant device type or model based on the initially selected devices. The primary reason for this automatic filtering behavior is that only one firmware image file for the category can be specified.

- NOTE: The update size limitation of either a single DUP or combined DUPs can be ignored if the Extended Storage feature is installed and enabled. For information on enabling extended storage, see Configuring CMC Extended Storage Card.
- **5.** Specify the firmware image file for the selected component(s) or devic(es). This is a Microsoft Windows Dell Update Package (DUP) file.
- 6. Select one of the following options:
 - **Reboot Now** Reboot immediately. The firmware update is applied immediately
 - On Next Reboot Manually reboot the server at a later time. The firmware update is applied after the next reboot.
 - NOTE: This step is not valid for Lifecycle Controller and 32-bit Diagnostics firmware update. A server reboot is not required for these devices.
- 7. Click Update. The firmware version is updated for the selected component or device.

Server Component Single Click Update Using Network Share

The Servers or server component update from a network share using Dell Repository Manager and Dell PowerEdge VRTX chassis integration simplifies the update by using customized bundle firmware, so that you can deploy faster and more easily. Update from a network share provides flexibility to update all the 12G server components at the same time with a single catalog either from a CIFS or from a NFS.

This method provides a quick and easy way to build a custom repository for connected systems that you own using the Dell Repository Manager and the chassis inventory file exported using the CMC Web interface. DRM enables you to create a fully customized repository that only includes the update packages for the specific system configuration. You can also build repositories that contain updates for only out-of-date devices, or a baseline repository that contains updates for all the devices. You can also create update bundles for Linux or Windows based on the update mode required. DRM enables you to save the repository to a CIFS or NFS share. The CMC Web interface enables you to configure the credentials and location details for the share. Using the CMC Web interface, you can then perform the server components update for a single server or multiple servers.

Pre-requisites for Using Network Share Update Mode

The following pre-requisites are required to update server component firmware using Network Share mode:

- The servers must belong to 12th or later generations and must have iDRAC Enterprise license.
- CMC version must be at version 2.0 or later.
- Lifecycle controller must be enabled on the servers.
- iDRAC Version 1.50.50 or later must be available on the 12th generation servers.
- Dell Repository Manager 1.8 or later must be installed on the system.
- You must have CMC Administrator privileges.

Upgrading Server Component Firmware From Network Share Using CMC Web Interface

To upgrade the server components firmware version to the next version using the **Update from Network Share** mode:

- In the CMC Web interface, in the system tree, go to Server Overview and then click Update > Server Component Update .
 - The Server Component Update page is displayed.
- 2. In the **Choose Update Type** section, select **Update from Network Share**. For more information, see Choosing Server Component Firmware Update Type.
- 3. If the Network Share is not connected, configure the Network Share for the chassis. To configure or edit the network share details, in the Network Share Properties table click **Edit**. For more information see Configuring Network Share Using CMC Web Interface.
- 4. Click Save Inventory Report to export the chassis inventory file that contains the components and firmware details. The Inventory.xml file is saved on an external system. The Dell Repository Manager uses the inventory.xml file to create customized bundles of updates. This Repository is stored in the CIFS or NFS Share configured by CMC. For information on creating a repository using the Dell Repository Manager see the Dell Repository Manager Data Center Version 1.8 User's Guide and the Dell Repository Manager Business Client Version 1.8 User's Guide available at dell.com/support/manuals.
- 5. Click Check for Updates to view the firmware updates available in the network share. The Component/Device Firmware Inventory section displays the current firmware versions of the components and devices across all the servers present in the chassis and firmware versions of the DUPs available in the Network Share.
 NOTE: Click Collapse against a slot to collapse the component and device firmware details for the specific slot. Alternatively, to view all the details again, click Expand.
- 6. In the **Component/Device Firmware Inventory** section, select the check box against **Select/Deselect All** to select all the supported servers. Alternatively, select the check box against the server for which you want to update the server component firmware. You cannot select individual components for the server.
- 7. Select one of the following options to specify if a system reboot is required after the updates are scheduled:
 - Reboot Now Updates are scheduled and the server is rebooted, immediately applying the updates to the server components.

- On Next Reboot Updates are scheduled but are applied only after the next server reboot.
- 8. Click **Update** to schedule firmware updates for the available components of the selected servers.

 A message is displayed based on the type of updates contained and asking you to confirm if you want to continue.
- 9. Click **OK** to continue and complete scheduling the firmware update for the selected servers. Note:
 - NOTE: The Job Status column displays the job status of the operations scheduled on the server. The job status is dynamically updated.

Supported Firmware Versions for Server Component Update

The following section provides the Server Component Update for CMC.

The following table lists the supported firmware versions for server components in a scenario where the existing CMC Firmware version is 3.0 and the server components are updated from N-1 version to N version.

NOTE: Server components firmware update from N-1 version to N version is successful when the CMC firmware is at version 2.0 or later, for all the 12th, 13th, and 14th generation servers mentioned in the following table.

Table 13. Supported Server Component Versions For Server Component Update to N version

Platform	Server Component	Previous Component Version (N-1 Version)	Updated Component Version (N Version)
M520	iDRAC	2.41.40.40	2.50.50.50
	Lifecycle Controller	2.41.40.40	2.50.50.50
	BIOS	2.3.2	2.4.2
M620	iDRAC	2.41.40.40	2.50.50.50
	Lifecycle Controller	2.41.40.40	2.50.50.50
	BIOS	2.5.2	2.4.2
M820	iDRAC	2.41.40.40	2.50.50.50
	Lifecycle Controller	2.41.40.40	2.50.50.50
	BIOS	2.3.2	2.3.3
M630	iDRAC	2.41.40.40	2.50.50.50
	Lifecycle Controller	2.41.40.40	2.50.50.50
	BIOS	2.4.2	2.5.4
M830	iDRAC	2.41.40.40	2.50.50.50
	Lifecycle Controller	2.41.40.40	2.50.50.50
	BIOS	2.4.2	2.5.4
M640	iDRAC	Not Applicable	3.10.10.10
	Lifecycle Controller	Not Applicable	3.10.10.10
	BIOS	Not Applicable	1.0.0

Deleting Scheduled Server Component Firmware Jobs

i NOTE: To use this feature, you must have an Enterprise License.

You can delete jobs scheduled for the selected components and/or devices across one or more servers.

Deleting Scheduled Server Component Firmware Jobs Using the Web Interface

To delete scheduled server component firmware jobs:

- 1. In the left pane, click Server Overview, and then click Update.
- 2. On the Server Component Update page, filter the component or device (optional).
- 3. In the **Job Status** column, if a check box is displayed next to the job status, it implies that a Lifecycle Controller job is in progress and currently in the indicated state. It can be selected for a job-deletion operation.
- 4. Click **Delete Job**. The jobs are deleted for the selected components or devices.

Updating Storage Component Using CMC Web Interface

Make sure the DUPs for the required storage components are downloaded.

To update the storage components:

- 1. In the left pane, click Chassis Overview > Storage > Update.
- On the Storage Component Update page, click Browse. The Choose to Upload File dialog box is displayed
- 3. Browse to location where the required DUP file was downloaded and saved from the Dell support Site and select the DUP file, and click **Open**.

The DUP file name and path are displayed in the Browse field.

4. Click Upload

The DUP is uploaded to CMC. The **Storage Component Update** section displays only the components that are supported by the downloaded DUP file. The current version, latest available version and the **Update** check box are displayed for the components.

- 5. Select the appropriate **Update** check boxes for the required components.
- 6. Click Update.

The firmware update action is initiated for the selected components. The progress is displayed in the **Update** column

After the action is complete, an appropriate message is displayed to indicate the completion or failure of the firmware update.

(i) NOTE:

- Servers must be turned off before updating the firmware.
- Component updates other corresponding components in the system similarly. For example The SPERC's updates similarly to the existing SPERCs and the EMMs updates similarly to the integrated EMMs.
- Click to view the HDD of different enclosures.

Recovering iDRAC Firmware Using CMC

iDRAC firmware is typically updated using iDRAC interfaces such as the iDRAC web interface, the SM-CLP command line interface, or operating system specific update packages downloaded from **support.dell.com**. For more information, see the *iDRAC User's Guide*.

Early generations of servers can have corrupted firmware recovered using the new update iDRAC firmware process. When CMC detects corrupted iDRAC firmware, it lists the server on the **Firmware Update** page. Complete the tasks mentioned in the Updating Server iDRAC Firmware.

Viewing Chassis Information and Monitoring Chassis and Component Health

You can view information and monitor the health of the following:

- Active and standby CMCs
- All severs and individual servers
- IO Module
- Fans
- Power Supply Units (PSUs)
- Temperature sensors
- Hard disk drives
- LCD assembly
- Storage controllers
- PCle devices

NOTE: The health of external components impacts the overall health of the storage component with existing storage health and integrated storage components in VRTX. It indicates that the external components will not impact the health of any component in the chassis.

Topics:

- · Viewing Chassis and Component Summaries
- Viewing Chassis Summary
- Viewing Chassis Controller Information and Status
- Viewing Information and Health Status of All Servers
- Viewing Health Status and Information for Individual Server
- Viewing Information and Health Status of the IOM
- Viewing Information and Health Status of Fans
- Viewing Front Panel Properties
- Viewing KVM Information and Health Status
- Viewing LCD Information and Health
- Viewing Information and Health Status of Temperature Sensors
- Viewing Storage Capacity and Status of the Storage Components

Viewing Chassis and Component Summaries

When you log in to the CMC web interface, the **Chassis Health** page displays the health of the chassis and its components. It displays a graphical view of the chassis and its components. It is dynamically updated, and the component sub-graphic overlays and text hints are automatically changed to reflect the current state.



To view the chassis health, click **Chassis Overview**. The system displays the overall health status of the chassis, active and standby CMCs, server modules, IO Module (IOM), fans, blowers, power supply units (PSUs), LCD assembly, storage controller, and PCle devices. Detailed information about each component is displayed when you click that component. In addition, the latest events in the CMC Hardware Log are also displayed. For more information, see the *Online Help*.

If your chassis is configured as a Group Lead, the **Group Health** page is displayed after login. It displays the chassis level information and alerts. All active, critical, and non-critical alerts are displayed.

Chassis Graphics

The chassis is represented by the front and back views (the upper and lower images respectively). Servers, DVDs, HDDs, KVMs, and LCD are shown in the front view and the remaining components are shown in the back view. Component selection is indicated by a blue cast and is controlled by clicking the image of the required component. When a component is present in the chassis, an icon of the component type is displayed in the graphics in the position (slot), where the component has been installed. Empty positions are shown with a charcoal gray background. The component icon visually indicates the state of the component. Other components display icons that visually represent the physical component. Pausing the cursor over a component displays a tool tip with additional information about that component.

Table 14. Server Icon States

Icon	Description	
	A server is present, turned on, and is operating normally.	
	A server is present, but turned off.	
	A server is present, but reporting a noncritical error.	

Table 14. Server Icon States (continued)

Icon	Description
×	A server is present, but reporting a critical error.
	A server is not present.

Selected Component Information

Information for the selected component is displayed in three independent sections:

- Health and Performance, and Properties Displays the active, critical, and non-critical events as displayed by the hardware logs and the performance data that vary with time.
- Properties Displays the component properties that do not vary with time, or that change only infrequently.
- Quick Links Provides links to navigate to the most frequently accessed pages, and also the most frequently performed
 actions. Only links applicable to the selected component are displayed in this section.

The following table lists the component properties and information displayed on the **Chassis Health** page in the web interface.

i NOTE: In Multi-Chassis Management (MCM), all the Quick Links associated with the servers are not displayed.

Table 15. Component properties

Component	Heath and Performance Properties	Properties	Quick Links
LCD Assembly	LCD Health	Chassis Power Button	Front Panel Configuration
	Chassis Health	Lock Control Panel LCD	
		LCD Language	
		LCD Orientation	
Active and	Redundancy Mode	Firmware	CMC Status
Standby CMCs	MAC Address	Standby Firmware	Networking
	• IPv4	Last Update	Firmware Update
	• IPv6	Hardware	
All Servers and	Power State	Name	Server Status
Individual Servers	Power Consumption	Model	Launch Remote Console
	Health	Service Tag	Launch iDRAC GUI
	Power Allocated	Host Name	Power Off Server
	Temperature	• iDRAC	Graceful Shutdown
		• CPLD	Remote File Share
		• BIOS	Deploy iDRAC Network
		• OS	Server Component Update

Table 15. Component properties (continued)

		 CPU Information Total System Memory 	(i) NOTE: Quick links for Power Off Server and Graceful Shutdown are displayed only if the server power state is On. If the server power state is Off, the quick link for Power On Server is displayed.
KVM Slot	Health	 KVM Mapped Slot 1: Front Panel USB/ Video Enabled Slot 2: Front Panel USB/ Video Enabled Slot 3: Front Panel USB/ Video Enabled Slot 4: Front Panel USB/ Video Enabled 	Front Panel Configuration
DVD Slot	HealthPower State	 DVD Mapped Slot 1: DVD Enabled Slot 2: DVD Enabled Slot 3: DVD Enabled Slot 4: DVD Enabled 	Front Panel Configuration
Disk Slot	HealthState	 Model Serial Number Power Status Firmware Version Size Type 	 Physical Disk Status Physical Disk Setup View Controller for this Physical Disk View Virtual Disks for this Physical Disk
Power Supply Units	Power Status	Capacity	Power Supply StatusPower ConsumptionSystem Budget
PCIe Devices	InstalledAssigned	 Model Server Slot Mapping Vendor ID Device ID Slot Type Allocated Power Fabric Power Status 	PCle StatusPCle Setup
Fans	• Speed	Warning Threshold	Fans Status

Table 15. Component properties (continued)

	PWM (% of Max)Fan Offset	Critical Threshold	• Fan Configuration
Blower	SpeedPWM (% of Max)	Warning ThresholdCritical Threshold	Fans StatusFan Configuration
	Enhanced Cooling Mode		
SPERC Slot	Installed	Model	Controller Status
	Assigned	Server Slot Mapping	Controller Setup
		Vendor ID	
		Device ID	
		Slot Type	
		Allocated Power	
		• Fabric	
		Power Status	
External Shared	Installed	Model	PCle Slot Status
PERC 8 card slot	Assigned	Server Slot Mapping	PCle Setup
		Vendor ID	
		Device ID	
		Slot Type	
		Allocated Power	
		• Fabric	
		Power Status	
IOM Slot	Power State	Model	IOM Status
	Role	Service Tag	Launch IOM GUI

Viewing Server Model Name and Service Tag

You can view the model name and service tag of each server instantly using the following steps:

- 1. In the left pane, under **Server Overview** tree node, all the servers (SLOT-01 to SLOT-04) appear in the servers list. If a server is not present in a slot, the corresponding image in the graphic is grayed out. When a full height server occupies slot 1 and slot 3, slot 3 will show the slot name as **Extension of 1**.
- 2. Pause the cursor over the slot name or slot number of a server. A tool tip is displayed with the server's model name and service tag (if available).

Viewing Chassis Summary

To view the chassis summary information, in the left pane, click **Chassis Overview** > **Properties** > **Summary**. The **Chassis Summary** page is displayed. For more information about this page, see the *Online Help*.

Viewing Chassis Controller Information and Status

To view the chassis controller information and status, in the CMC Web interface, click **Chassis Overview > Chassis Controller**.

The Chassis Controller Status page is displayed. For more information, see the Online Help.

Viewing Information and Health Status of All Servers

To view the health status of all the servers, do one of the following:

- Click **Chassis Overview**. The **Chassis Health** page displays a graphical overview of all the servers installed in the chassis. Server health status is indicated by the overlay of the server subgraphic. For more information about the chassis health, see the *Online Help*.
- Click **Chassis Overview** > **Server Overview**. The **Servers Status** page provides an overview of the servers in the chassis. For more information, see the *Online Help*.

Viewing Health Status and Information for Individual Server

To view health status for individual servers, do any of the following:

- 1. Go to Chassis Overview > Properties > Health.
 - The **Chassis Health** page displays a graphical overview of all the servers installed in the chassis. Server health status is indicated by the overlay of the server subgraphic. Move the cursor to hover over an individual server subgraphic. A corresponding text hint or screen tip provides additional information for that server. Click the server subgraphic to view the I/O Module information on the right. For more information, see the *Online Help*.
- 2. Go to Chassis Overview and expand Server Overview in the left pane. All the servers (1-4) appear in the expanded list. Click the server (slot) you want to view.
 - The **Server Status** page (separate from the **Servers Status** page) provides the health status of the server in the chassis and a launch point to the iDRAC web interface, which is the firmware used to manage the server. For more information, see the *Online Help*.
 - NOTE: To use the iDRAC web interface, you must have an iDRAC user name and password. For more information about iDRAC and the using the iDRAC web interface, see the *Integrated Dell Remote Access Controller User's Guide*.

Viewing Information and Health Status of the IOM

To view health status of the IOMs, in the CMC Web interface, do any of the following:

- 1. Click Chassis Overview.
 - The **Chassis Health** page is displayed. The graphics in the left pane displays the rear, front, and side view of the chassis and contains the health status for the IOM. IOM health status is indicated by the overlay of the IOM sub-graphic. Move the cursor over the individual IOM sub-graphic. The text hint provides additional information about that IOM. Click the IOM sub-graphic to view the IOM information in the right pane.
- Go to Chassis Overview > I/O Module Overview.
 The I/O Module Status page provides an overview of IOM associated with the chassis. For more information, see the Online Help.
- NOTE: After updating or power cycling the IOM/IOA, make sure that the operating system of the IOM/IOA is also booted correctly. Else, the IOM status is displayed as "Offline".

Viewing Information and Health Status of Fans

CMC controls the speed of the chassis fan by increasing or decreasing the fan speed on the basis of system events. You can run the fan in three modes such as Low, Medium, and High. For more information about configuring a fan, see the *Online Help*.

To set up the properties of fans by using RACADM commands, type the following command at the CLI interface.

```
racadm fanoffset [-s <off|low|medium|high>]
```

For more information about the RACADM commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/cmcmanuals**.

NOTE: The CMC monitors the temperature sensors in the chassis and automatically adjust the fan speed as needed. However, you can override to maintain a minimum fan speed by the racadm fanoffset command. When overriden using this command, the CMC will always run the fan to the selected speed even though the chassis does not require the fans to run at that speed.

CMC generates an alert and increases the fan speeds when the following events occur:

- CMC ambient temperature threshold is exceeded.
- A fan stops functioning.
- A fan is removed from the chassis.
- NOTE: During updates of CMC or iDRAC firmware on a server, some or all of the fan units in the chassis rotates at 100%. This is normal.

To view the health status of fans, in the CMC Web interface, do any of the following:

1. Go to Chassis Overview.

The **Chassis Health** page is displayed. The lower section of chassis graphics provides the left view of the chassis and contains the health status of the fans. Fan health status is indicated by the overlay of the fan sub-graphic. Move the cursor over the fan sub-graphic. The text hint provides additional information about a fan. Click the fan subgraphic to view the fan information in the right pane.

2. Go to Chassis Overview > Fans.

The **Fans Status** page provides the status, speed measurements in revolutions per minute (RPMs), and threshold values of the fans in the chassis. There can be one or more fans.

- NOTE: In the event of a communication failure between CMC and the fan unit, CMC cannot obtain or display the health status for the fan unit.
- NOTE: The following message is displayed when both the fans are not present in the slots or if a fan is rotating at a low speed:

```
Fan <number> is less than the lower critical threshold.
```

For more information, see the Online Help.

Configuring Fans

Fan Offset — A feature to provide increased cooling to the storage and PCle regions of the chassis. This feature allows you to increase the airflow delivery to the HDDs, Shared PERC controllers, and PCle card slots. An example usage of the Fan Offset is when you use high-power or custom PCle cards that require more cooling than normal. The Fan Offset feature has options of Off, Low, Medium, and High. These settings correspond to a fan speed offset (increase) of 20%, 50%, and 100% of the maximum speed respectively. There are also minimum speeds setup for each option, which are 35% for Low, 65% for Medium, and 100% for High.

Using the Medium Fan Offset setting for example, increases the speed of fans 1–6 by 50% of its maximum speed. The increase is above the speed already set by the system for cooling on the basis of installed hardware configuration.

With any of the Fan Offset options enabled, the power consumption will be increased. The system will be louder with the Low offset, noticeably louder with the Medium offset, and significantly louder with the High offset. When the Fan Offset option is not enabled, the fan speeds will be reduced to the default speeds required for system cooling for the installed hardware configuration.

To set the offset feature, go to **Chassis Overview** > **Fans** > **Setup**. On the **Advanced Fan Configurations** page, in the **Fan Configuration** table, from the **Value** drop-down menu corresponding to **Fan Offset**, select an option appropriately.

For more information about the Fan Offset feature, see the Online Help.

For setting up these features by using RACADM commands, user the following command:

```
racadm fanoffset [-s <off|low|medium|high>]
```

For more information about the fan offset-related RACADM commands, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/Manuals.

Enhanced Cooling Mode (ECM) — Is a feature in the CMC that allows increased cooling capacity for the servers installed within the PowerEdge VRTX chassis. Example uses for ECM are operation in a high ambient environment or using servers with high power (≥120W) CPUs installed. The increased cooling capacity is achieved by allowing the four chassis blower modules to run at a higher speed. As a result, the system power consumption and noise level may be increased when ECM is enabled.

When enabled, ECM will only increase the cooling capacity to the server slots within the chassis. It is also important to note that ECM is not designed to provide increased cooling to the servers at all times. Even with ECM enabled, the higher blower speeds will only be seen when the increased cooling is needed. Examples of this situation include high levels of server utilization or stress, and high ambient temperatures.

By default, ECM is off. When ECM is enabled, the blowers have the capability to deliver approximately 20% more airflow per blade.

To set the ECM mode, go to **Chassis Overview** > **Fans** > **Setup**. On the **Advanced Fan Configurations** page, in the **Blower Configuration**table, from the **Value** drop-down menu corresponding to **Enhanced Cooling Mode**, select an option appropriately.

For more information about the ECM feature, see the Online Help.

Viewing Front Panel Properties

To view the front panel properties:

- 1. In the left pane, click Chassis Overview > Front Panel.
- 2. On the **Properties** page, you can view the following:
 - Power Button Properties
 - LCD Properties
 - KVM Properties
 - DVD Drive Properties

Viewing KVM Information and Health Status

To view the health status of the KVMs associated with the chassis, do any of the following:

- 1. Click Chassis Overview.
 - The **Chassis Health** page is displayed. The left pane displays the front view of the chassis and contains the health status of a KVM. KVM health status is indicated by the overlay of the KVM sub-graphic. Move the pointer over an KVM sub-graphic and a corresponding text hint or screen tip is displayed. The text hint provides additional information about the KVM. Click the KVM sub-graphic to view the KVM information in the right pane.
- Alternatively, click Chassis Overview > Front Panel.
 On the Status page, under the KVM Properties section, you can view the status and properties of a KVM associated with the chassis. For more information, see the Online Help.

Viewing LCD Information and Health

To view the health status of an LCD:

- 1. In the left pane, click Chassis Overview.
 - The **Chassis Health** page is displayed. The left pane displays the front view of the chassis. LCD health status is indicated by the overlay of the LCD sub-graphic.
- 2. Move the cursor over the LCD subgraphic. The corresponding text hint or screen tip provides additional information on the LCD.
- 3. Click the LCD sub-graphic to view the LCD information in the right pane. For more information, see the Online Help.
 Alternatively, go to Chassis Overview > Front Panel > Properties > Status. On the Status page, under the LCD Properties, you can view the status of the LCD available on the chassis. For more information, see Online Help.

Viewing Information and Health Status of Temperature Sensors

To view the health status of the temperature sensors:

In the left pane, click Chassis Overview > Temperature Sensors.

The **Temperature Sensors Status** page displays the status and readings of the temperature probes on the entire chassis (chassis and servers). For more information, see *Online Help*.

NOTE: The temperature probes value cannot be edited. Any change beyond the threshold generates an alert that causes the fan speed to vary. For example, if the CMC ambient temperature probe exceeds the threshold, the speed of the fans on the chassis increases.

Viewing Storage Capacity and Status of the Storage Components

To view the capacity and fault-tolerant status of the storage components, do one of the following:

1. Go to Chassis Overview

The **Chassis Health** page is displayed. The Storage capacity details, the Fault Tolerant Mode (Active/Passive), and Fault Tolerant Status (Enabled) information is displayed on the right pane. This fault-tolerance information is displayed only if the fault tolerant feature is enabled for the storage components.

The lower section of chassis graphics provides the left view of the chassis. Move the cursor over the storage component sub-graphic. The text hint provides additional information about the storage component. Click the storage component subgraphic to view the related information in the right pane.

2. Alternatively, in the left pane, click Chassis Overview > Storage > Properties > Status.

The **Storage Overview** page is displayed with the following information:

- View the graphic summary of the physical disk drives installed in the chassis and their status.
- View the summary of all the storage components with links to their respective pages.
- View the used capacity and total capacity of the storage.
- View controller information.
 - (Integrated 2). NOTE: In case of a fault-tolerant controller, the name format is: Shared < PERC number > (Integrated < number >). For example, the active controller is Shared PERC8 (Integrated 1) and the peer controller is Shared PERC8 (Integrated 2).
- View recently-logged storage events.
- (i) NOTE: For more information, see the Online Help.

Configuring CMC

Chassis Management Controller enables you to configure properties, set up users, and alerts to perform remote management tasks.

Before you begin configuring the CMC, you must first configure the CMC network settings to allow CMC to be managed remotely. This initial configuration assigns the TCP/IP networking parameters that enable access to the CMC. For more information, see Setting Up Initial Access to CMC.

You can configure CMC using Web interface or RACADM.

NOTE: When you configure CMC for the first time, you must be logged in as root user to execute RACADM commands on a remote system. Another user can be created with privileges to configure CMC.

After setting up the CMC and performing the basic configurations, you can do the following:

- Modify the network settings, if required.
- Configure interfaces to access CMC.
- Configure LCD display.
- Set up chassis groups, if required.
- Configure servers, I/O module, or front panel.
- Configure VLAN settings.
- Obtain the required certificates.
- Add and configure CMC users with privileges.
- Configure and enable e-mail alerts and SNMP traps.
- Set the power cap policy, if required.
- NOTE: The following characters cannot be used in the property strings of both the CMC interfaces (GUI and CLI):
 - &#
 - < and > together
 - ; (semicolon)

Topics:

- Viewing and Modifying CMC Network LAN Settings
- Configuring CMC Network and Login Security Settings
- Configuring Virtual LAN Tag Properties for CMC
- Federal Information Processing Standards
- Configuring Services
- Configuring CMC Extended Storage Card
- Setting Up Chassis Group
- Chassis Configuration Profiles
- Configuring Multiple CMCs Using RACADM
- Configuring Multiple CMCs through RACADM Using Chassis Configuration Profiles
- Viewing and Ending CMC Sessions

Viewing and Modifying CMC Network LAN Settings

The LAN settings, such as community string and SMTP server IP address, affect both the CMC and the external settings of the chassis.

If you have two CMCs (active and standby) on the chassis that are connected to the network, the standby CMC automatically assumes the network settings of the active CMC in the event of failover.

When IPv6 is enabled at boot time, three router solicitations are sent after every four seconds. If external network switches are running the Spanning Tree Protocol (STP), the external switch ports may be blocked for more than 12 seconds in which

the IPv6 router solicitations are sent. In such cases, there may be a period when IPv6 connectivity is limited, until router advertisements are gratuitously sent by the IPv6 routers.

- i NOTE: Changing the CMC network settings may disconnect your current network connection.
- i NOTE: You must have Chassis Configuration Administrator privilege to set up CMC network settings.

Viewing and Modifying CMC Network LAN Settings Using CMC Web Interface

To view and modify the CMC LAN network settings using CMC Web interface:

- 1. In the left pane, click **Chassis Overview**, and then click **Network**. The **Network Configuration** page displays the current network settings.
- 2. Modify the general, IPv4, or IPv6 settings as required. For more information, see the Online Help.
- 3. Click Apply Changes for each section to apply the settings.

Viewing and Modifying CMC Network LAN Settings Using RACADM

To view IPv4 settings, use the objects from the **cfgCurrentLanNetworking** group with the following getniccfg and getconfig subcommands.

To view IPv6 settings, use the objects from the cfglpv6LanNetworking group with the getconfig subcommand.

To view IPv4 and IPv6 addressing information for the chassis, use the getsysinfo subcommand.

For more information about the subcommands and objects, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Enabling the CMC Network Interface

To enable or disable the CMC network interface for both IPv4 and IPv6, type:

```
racadm config -g cfgLanNetworking -o cfgNicEnable 1
racadm config -g cfgLanNetworking -o cfgNicEnable 0
```

i NOTE: The CMC NIC is enabled by default.

To enable or disable the CMC IPv4 addressing, type:

```
racadm config -g cfgLanNetworking -o cfgNicIPv4Enable

1
racadm config -g cfgLanNetworking -o cfgNicIPv4Enable

0
```

i NOTE: The CMC IPv4 addressing is enabled by default.

To enable or disable the CMC IPv6 addressing, type:

```
racadm config -g cfgIpv6LanNetworking -o cfgIPv6Enable

1
racadm config -g cfgIpv6LanNetworking -o cfgIPv6Enable

0
```

- (i) NOTE: Note the following:
 - There is a 30 seconds delay between changing a network setting and actually applying it.
 - The CMC IPv6 addressing is disabled by default.

By default, for IPv4, the CMC requests and automatically obtains a CMC IP address from the Dynamic Host Configuration Protocol (DHCP) server. You can disable the DHCP feature and specify static CMC IP address, gateway, and subnet mask.

For an IPv4 network, to disable DHCP and specify static CMC IP address, gateway, and subnet mask, type:

```
racadm config -g cfgLanNetworking -o cfgNicUseDHCP 0
racadm config -g cfgLanNetworking -o cfgNicIpAddress <static IP address>
racadm config -g cfgLanNetworking -o cfgNicGateway <static gateway>
racadm config -g cfgLanNetworking -o cfgNicNetmask <static subnet mask>
```

By default, for IPv6, the CMC requests and automatically obtains a CMC IP address from the IPv6 autoconfiguration mechanism.

For an IPv6 network, to disable the Autoconfiguration feature and specify a static CMC IPv6 address, gateway, and prefix length, type:

```
racadm config -g cfgIPv6LanNetworking -o cfgIPv6AutoConfig 0 racadm config -g cfgIPv6LanNetworking -o cfgIPv6Address <IPv6 address> racadm config -g cfgIPv6LanNetworking -o cfgIPv6PrefixLength 64 racadm config -g cfgIPv6LanNetworking -o cfgIPv6Gateway <IPv6 address>
```

Enabling or Disabling DHCP for the CMC Network Interface Address

When enabled, the CMC's DHCP for NIC address feature requests and obtains an IP address from the Dynamic Host Configuration Protocol (DHCP) server automatically. This feature is enabled by default.

You can disable the DHCP for NIC address feature and specify a static IP address, subnet mask, and gateway. For more information, see Setting Up Initial Access to CMC.

Enabling or Disabling DHCP for DNS IP Addresses

By default, the CMC's DHCP for DNS address feature is disabled. When enabled, this feature obtains the primary and secondary DNS server addresses from the DHCP server. While using this feature, you do not have to configure static DNS server IP addresses.

To disable the DHCP for DNS address feature and specify the static preferred and alternate DNS server addresses, type:

```
racadm config -g cfgLanNetworking -o cfgDNSServersFromDHCP 0
```

To disable the DHCP for DNS address feature for IPv6 and specify the static preferred and alternate DNS server addresses, type:

```
racadm config -g cfgIPv6LanNetworking -o cfgIPv6DNSServersFromDHCP6 0
```

Setting Static DNS IP addresses

i) NOTE: The static DNS IP addresses settings are not valid unless the DCHP for DNS address feature is disabled.

For IPv4, to set the preferred primary and secondary DNS IP server addresses, type:

```
racadm config -g cfgLanNetworking -o cfgDNSServer1 <IP-address>
racadm config -g cfgLanNetworking -o cfgDNSServer2 <IPv4-address>
```

For IPv6, to set the preferred and secondary DNS IP Server addresses, type:

```
racadm config -g cfgIPv6LanNetworking -o cfgIPv6DNSServer1 <IPv6-address>
```

```
racadm config -g cfgIPv6LanNetworking -o
cfgIPv6DNSServer2 <IPv6-address>
```

Configuring IPv4 and IPv6 DNS Settings

• **CMC Registration** — To register the CMC on the DNS server, type:

```
racadm config -g cfgLanNetworking -o cfgDNSRegisterRac 1
```

- NOTE: Some DNS servers only register names of 31 characters or fewer. Make sure the designated name is within the DNS required limit.
- NOTE: The following settings are valid only if you have registered the CMC on the DNS server by setting cfqDNSRegisterRac to 1.
- **CMC Name** By default, the CMC name on the DNS server is cmc-<*service tag>*. To change the CMC name on the DNS server, type:

```
racadm config -g cfgLanNetworking -o cfgDNSRacName < name>
```

where <name> is a string of up to 63 alphanumeric characters and hyphens. For example: cmc-1, d-345.

- NOTE: If a DNS Domain name is not specified then the maximum number of characters is 63. If a domain name is specified then the number of characters in CMC name plus the number of characters in the DNS Domain Name must be less than or equal to 63 characters.
- DNS Domain Name The default DNS domain name is a single blank character. To set a DNS domain name, type:

```
racadm config -g cfgLanNetworking -o
cfgDNSDomainName <name>
```

where <name> is a string of up to 254 alphanumeric characters and hyphens. For example: p45, a-tz-1, r-id-001.

Configuring Auto Negotiation, Duplex Mode, and Network Speed for IPv4 and IPv6

When enabled, the auto negotiation feature determines whether the CMC automatically sets the duplex mode and network speed by communicating with the nearest router or switch. By default, auto negotiation feature is enabled.

You can disable auto negotiation and specify the duplex mode and network speed by typing:

```
racadm config -g cfgNetTuning -o cfgNetTuningNicAutoneg 0
racadm config -g cfgNetTuning -o cfgNetTuningNicFullDuplex <duplex mode>
```

where:

<duplex mode> is 0 (half duplex) or 1 (full duplex, default)

```
racadm config -g cfgNetTuning -o cfgNetTuningNicSpeed < speed>
```

where:

<speed> is 10 or 100 (default).

Setting the Maximum Transmission Unit for IPv4 and IPv6

The Maximum Transmission Unit (MTU) property allows you to set a limit for the largest packet that can be passed through the interface. To set the MTU, type:

```
racadm config -g cfgNetTuning -o cfgNetTuningMtu <mtu>
```

where $\langle mtu \rangle$ is a value between 576–1500 (inclusive; default is 1500).

NOTE: IPv6 requires a minimum MTU of 1280. If IPv6 is enabled, and cfgNetTuningMtu is set to a lower value, the CMC uses an MTU of 1280.

Configuring CMC Network and Login Security Settings

The IP address blocking and User blocking features in CMC allow you to prevent security issues due to password guessing attempts. This feature enables you to block a range of IP addresses and users who can access CMC. By deafult, the IP address blocking feature is enabled in CMC. You can set the IP range attributes using CMC web interface or RACADM. To use the IP address blocking and user blocking features, enable the options using CMC web interface or RACADM. Configure the login lockout policy settings to enable you to set the number of unsuccessful login attempts for a specific user or for an IP address. After exceeding this limit, the blocked user can log in only after the penalty time expires.

i NOTE: Blocking by IP address is applicable only for IPV4 addresses.

Configuring IP Range Attributes Using CMC Web Interface

(i) NOTE: To perform the following task, you must have Chassis Configuration Administrator privilege.

To configure the IP range attributes using CMC web interface:

- 1. In the left pane, go to Chassis Overview and click Network > Network. The Network Configuration page is displayed.
- 2. In the IPv4 Settings section, click **Advanced Settings**.

The Log in Security page is displayed.

Alternatively, to access the Log in Security page, in the left pane, go to Chassis Overview, click Security > Log in.

- To enable the IP range checking feature, in the IP Range section, select the IP Range Enabled option.
 The IP Range Address and IP Range Mask fields are activated.
- 4. In the IP Range Address and IP Range Mask fields, type the range of IP addresses and IP range masks that you want to block from accessing CMC.

For more information, see the Online Help.

5. Click Apply to save your settings.

Configuring IP Range Attributes Using RACADM

You can configure the following IP Range attributes for CMC using RACADM:

- IP range checking feature
- Range of IP addresses that you want to block from accessing CMC
- IP Range Mask that you want to block from accessing CMC

IP filtering compares the IP address of an incoming login to the IP address range that is specified. A login from the incoming IP address is allowed only if both the following are identical:

- cfgRacTunelpRangeMask bit-wise and with incoming IP address
- cfgRacTunelpRangeMask bit-wise and with cfgRacTunelpRangeAddr
- To enable the IP range checking feature, use the following property under cfgRacTuning group:

cfgRacTuneIpRangeEnable <0/1>

 To specify the range of IP addresses that you want to block from accessing CMC, use the following property under cfgRacTuning group:

cfgRacTuneIpRangeAddr

• To specify the IP Range Mask that you want to block from accessing CMC, use the following property under cfgRacTuning group:

cfgRacTuneIpRangeMask

Configuring Virtual LAN Tag Properties for CMC

Virtual LANfunctionality enables multiple VLANs to coexist on the same physical network cable and to segregate the network traffic for security or load management purposes. When you enable the VLAN functionality, each network packet is assigned a VLAN tag.

Configuring Virtual LAN Tag Properties for CMC Using RACADM

1. Enable the Virtual LAN (VLAN) capabilities of the external chassis management network:

```
racadm config -g cfgLanNetworking -o cfgNicVLanEnable 1
```

2. Specify the VLAN ID for the external chassis management network:

```
racadm config -g cfgLanNetworking -o cfgNicVlanID <VLAN id>
```

The valid values for <VLAN id> are 1-4000 and 4021-4094. Default value is 1.

For example:

```
racadm config -g cfgLanNetworking -o cfgNicVlanID
```

3. Then, specify the VLAN priority for the external chassis management network:

```
racadm config -g cfgLanNetworking -o
cfgNicVLanPriority <VLAN priority>
```

The valid values for $\langle VLAN \ priority \rangle$ are 0-7. Default value is 0.

For example:

```
racadm config -g cfgLanNetworking -o cfgNicVLanPriority 7
```

You can also specify both the VLAN ID and the VLAN priority with a single command:

```
racadm setniccfg -v <VLAN id> <VLAN priority>
```

For example:

```
racadm setniccfg -v 1 7
```

4. To remove the CMC VLAN, disable the VLAN capabilities of the external chassis management network:

```
racadm config -g cfgLanNetworking -o
cfgNicVLanEnable 0
```

You can also remove the CMC VLAN using the following command:

```
racadm setniccfg -v
```

Configuring Virtual LAN Tag Properties for CMC Using Web Interface

To configure Virtual LAN(VLAN) for CMC using the CMC Web interface:

- 1. Go to any of the following pages:
 - In the left pane, click Chassis Overview and click Network > VLAN.
 - In the left pane, click Chassis Overview > Server Overview and click Network > VLAN.

The **VLAN Tag Settings** page is displayed. VLAN tags are chassis properties. They remain with the chassis even when a component is removed.

- 2. In the **CMC** section, enable VLAN for CMC, set the priority and assign the ID. For more information about the fields, see the *Online Help*.
- 3. Click Apply. The VLAN tag settings are saved.

You can also access this page from the Chassis Overview > Servers > Setup > VLAN.

Federal Information Processing Standards

The agencies and contractors of the Federal government of the United States use Federal Information Processing Standards (FIPS), a computer security standard, which is related to all applications that have communicative interfaces. The 140–2 comprises of four levels — Level 1, Level 2, Level 3, and Level 4. The FIPS 140–2 series stipulate that all communicative interfaces must have the following security properties:

- authentication
- confidentiality
- message integrity
- non-repudiation
- availability
- access control

If any of the properties depend on cryptographic algorithms, then FIPS must approve these algorithms.

By default, the FIPS mode is disabled. When FIPS is enabled, the minimum key size for OpenSSL FIPS is SSH-2 RSA 2048 bits.

i NOTE: PSU firmware update is not supported when the FIPS mode is enabled in the chassis.

For more information, see CMC Online Help.

The following features/applications support FIPS.

- Web GUI
- RACADM
- WSMan
- SSH v2
- SMTP
- Kerberos
- NTP Client
- NFS

NOTE: SNMP is not FIPS-compliant. In FIPS mode, all SNMP features except Message Digest algorithm version 5 (MD5) authentication work.

Enabling FIPS Mode Using CMC Web Interface

To enable FIPS:

- In the left pane, click Chassis Overview.
 The Chassis Health page is displayed.
- 2. On the menu bar, click Network.

The **Network Configuration** page is displayed.

 Under the Federal Information Processing Standards (FIPS) section, from the FIPS Mode drop-down menu, select Enabled.

A message is displayed that enabling FIPS resets CMC to the default settings.

4. Click **OK** to proceed.

Enabling FIPS Mode Using RACADM

To enable FIPS mode, run the following command:

racadm config -q cfqRacTuning -o cfqRacTuneFipsModeEnable 1

Disabling FIPS Mode

To disable FIPS mode, reset CMC to the default factory settings.

Configuring Services

You can configure and enable the following services on CMC:

- CMC serial console Enable access to CMC using the serial console.
- Web Server Enable access to CMC web interface. Disabling the web server also disables Remote RACADM.
- SSH Enable access to CMC through firmware RACADM.
- Telnet Enable access to CMC through firmware RACADM
- RACADM Enable access to CMC using RACADM.
- SNMP Enable CMC to send SNMP traps for events.
- Remote Syslog Enable CMC to log events to a remote server. To use this feature, you must have an Enterprise license.

CMC includes a web server that is configured to use the industry-standard SSL security protocol to accept and transfer encrypted data from and to clients over the Internet. The web server includes a Dell self-signed SSL Digital Certificate (Server ID), and is responsible for accepting and responding to secure HTTP requests from clients. This service is required by the web interface and remote RACADM CLI tool for communicating with CMC.

If the web server resets, wait at least one minute for the services to become available again. A web server reset usually happens as a result of any of the following events:

- Network configuration or network security properties are changed through the CMC web user interface or RACADM.
- Web server port configuration is changed through the web user interface or RACADM.
- CMC is reset.
- A new SSL server certificate is uploaded.
- i NOTE: To modify service settings, you must have the Chassis Configuration Administrator privilege.

Remote syslog is an additional log target for CMC. After you configure the remote syslog, each new log entry generated by CMC is forwarded to the respective destinations.

NOTE: Because the network transport for the forwarded log entries is UDP, there is no guaranteed delivery of log entries, nor is there any feedback to CMC about whether the log entries were received successfully.

Configuring Services Using CMC Web Interface

To configure CMC services using CMC web interface:

- 1. In the left pane, click **Chassis Overview**, and then click **Network** > **Services**. The **Services Management** page is displayed.
- 2. Configure the following services as required:
 - CMC Serial
 - Web Server
 - SSH

- Telnet
- Remote RACADM
- SNMP
- Remote Syslog

For information about the fields, see the Online Help.

3. Click Apply, and then update all default time-out and maximum time-out limits.

Configuring Services Using RACADM

To enable and configure the various services, use the following RACADM objects:

- cfgRacTuning
- cfgRacTuneRemoteRacadmEnable

For more information about these objects, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

If the firmware on the server does not support a feature, configuring a property related to that feature displays an error. For example, using RACADM to enable remote syslog on an unsupported iDRAC displays an error message.

Similarly, when displaying the iDRAC properties using the RACADM getconfig command, the property values are displayed as N/A for an unsupported feature on the server.

For example:

```
$ racadm getconfig -g cfgSessionManagement -m server-1
# cfgSsnMgtWebServerMaxSessions=N/A
# cfgSsnMgtWebServerActiveSessions=N/A
# cfgSsnMgtWebServerTimeout=N/A
# cfgSsnMgtSSHMaxSessions=N/A
# cfgSsnMgtSSHActiveSessions=N/A
# cfgSsnMgtSSHTimeout=N/A
# cfgSsnMgtSSHTimeout=N/A
# cfgSsnMgtTelnetMaxSessions=N/A
# cfgSsnMgtTelnetActiveSessions=N/A
# cfgSsnMgtTelnetActiveSessions=N/A
# cfgSsnMgtTelnetTimeout=N/A
```

Configuring CMC Extended Storage Card

You can enable or repair the optional Removable Flash Media for use as an extended non-volatile storage. Some CMC features depend on extended nonvolatile storage for their operation.

To enable or repair the Removable Flash Media using the CMC web interface:

- 1. In the left pane, go to Chassis Overview, and then click Chassis Controller > Flash Media.
- 2. On the Removable Flash Media page, from the drop-down menu, select one of the following as appropriate:
 - Repair active controller media
 - Stop using flash media for storing chassis data

For more information about these options, see the Online Help.

3. Click **Apply** to apply the selected option.

If two CMCs are present in the chassis, both CMCs (active and standby) must contain flash media; else, the Extended Storage functionality will be degraded unless both the active and standby CMCs contain flash media.

Setting Up Chassis Group

CMC enables you to monitor multiple chassis from a single lead chassis. When a chassis group is enabled, CMC in the lead chassis generates a graphical display of the status of the lead chassis and all member chassis within the chassis group. To use this feature, you must have an Enterprise License.

The Chassis group features are:

• Displays images portraying the front and back of each chassis, a set for the leader and a set for each member.

- Health concerns for the leader and members of a group are recognized by red or yellow overlays and an X or an ! on the component with the symptoms. Details are visible below the chassis image when you click the chassis image or **Details**.
- Quick launch links are available for opening member chassis's or server's web pages.
- A server and Input/Output inventory is available for a group.
- A selectable option is available to synchronize a new member's properties to the leader's properties when the new member is added to the group.

A chassis group may contain a maximum of eight members. Also, a leader or member can only participate in one group. You cannot join a chassis, either as a leader or member, that is part of a group to another group. You can delete the chassis from a group and add it later to a different group.

To set up the Chassis Group using the CMC web interface:

- 1. Log in with chassis administrator privileges to the leader chassis.
- 2. Click Setup > Group Administration.
- 3. On the Chassis Group page, under Role, select Leader. A field to add the group name is displayed.
- 4. Type the group name in the Group Name field, and then click Apply.
 - i NOTE: The same rules that apply for a domain name apply to the group name.

When the chassis group is created, the GUI automatically switches to the **Chassis Group** page. The left pane indicates the group by the group name and the lead chassis, and the unpopulated member chassis appear in the left pane.

Adding Members To Chassis Group

After the Chassis Group is set up, to add members to the group:

- 1. Log in with chassis administrator privileges to the leader chassis.
- 2. Select the lead chassis in the tree.
- 3. Click Setup > Group Administration.
- 4. Under Group Management, enter the member's IP address or DNS name in the Hostname/IP Address field.
- 5. In the User Name field, enter a user name with chassis administrator privileges for the member chassis.
- **6.** Type the corresponding password in the **Password** field.
- 7. Optionally, select Sync New Member with Leader Properties to push leader properties to the member.
- 8. Click Apply.
- 9. To add a maximum of eight members, complete the tasks in step 4 through step 8. The chassis names of the new members appear in the **Members** dialog box.
 - NOTE: The credentials entered for a member are passed securely to the member chassis to establish a trust relationship between the member and lead chassis. The credentials are not persisted on either chassis, and are never exchanged again after the initial trust relationship is established.

Removing a Member from the Leader

You can remove a member from the group from the lead chassis. To remove a member:

- 1. Log in with chassis administrator privileges to the leader chassis.
- 2. In the left pane, select the lead chassis.
- 3. Click Setup > Group Administration.
- 4. From the Remove Members list, select the member's name to be deleted, and then click Apply.

The lead chassis then communicates to the member or members, if more than one is selected, that it has been removed from the group. The member name is removed. The member chassis may not receive the message, if a network issue prevents contact between the leader and the member. In this case, disable the member from the member chassis to complete the removal.

Disbanding a Chassis Group

To disband a chassis group from the lead chassis:

- 1. Log in with administrator privileges to the leader chassis.
- 2. Select the lead chassis in the left pane.
- 3. Click Setup > Group Administration.
- 4. In the Chassis Group page, under Role, select None, and then click Apply.

The lead chassis then communicates to all the members that they have been removed from the group. The lead chassis can be assigned as a leader or member of a new group.

If a network issue prevents contact between the leader and the member, the member chassis may not receive the message. In this case, disable the member from the member chassis to complete the removal process.

Disabling an Individual Member at the Member Chassis

Sometimes a member cannot be removed from a group by the lead chassis. This can happen if network connectivity to the member is lost. To remove a member from a group at the member chassis:

- 1. Log in with chassis administrator privileges to the member chassis.
- 2. In the left pane, click Chassis Overview > Setup > Group Administration.
- 3. Select None, and then click Apply.

Accessing the Web page of a Member Chassis or Server

You can access the web page of the member chassis, remote console of the server, or the web page of the iDRAC server from the lead chassis group page. If the member device has the same login credentials as the lead chassis, you can use the same credentials to access the member device.

NOTE: Single Sign-On and Smart Card Login are not supported in Multiple Chassis Management. To access members by Single Sign On from lead chassis requires a common username or password between Lead and members. Use of common username or password works only with Active Directory, local, and LDAP users.

To navigate to member devices:

- 1. Log in to the lead chassis.
- 2. Select Group: name in the tree.
- 3. If a member CMC is the required destination, select Launch CMC for the required chassis.

If a server in a chassis is the required destination:

- a. Select the image of the destination chassis.
- b. In the chassis image that is displayed in the **Health** section, select the server.
- **c.** In the box labeled **Quick Links**, select the destination device. A new window is displayed with the destination page or login screen.
 - i NOTE: In MCM, all the Quick Links associated with the servers are not displayed.

Propagating Leader Chassis Properties to Member Chassis

You can apply the properties from the leader to the member chassis of a group. To synchronize a member with the leader properties:

- 1. Login with administrator privileges to the leader chassis.
- 2. Select the Lead chassis in the tree.
- 3. Click Setup > Group Administration.
- 4. In the Chassis Properties Propagation section, select one of the propagation types:
 - On-Change Propagation Select this option for automatic propagation of the selected chassis property settings. The property changes are propagated to all current group members, whenever lead properties are changed.
 - Manual Propagation Select this option for manual propagation of the chassis group leader properties with its
 members. The lead chassis property settings are propagated to group members only when a lead chassis administrator
 clicks **Propagate**.

- In the Propagation Properties section, select the categories of lead configuration properties to be propagated to member chassis.
 - Select only those setting categories that you want identically configured, across all members of the chassis group. For example, select **Logging and Alerting Properties** category, to enable all chassis in the group to share the logging and alerting configuration settings of the lead chassis.
- 6. Click Save.

If **On-Change Propagation** is selected, the member chassis take on the properties of the leader. If **Manual Propagation** is selected, click **Propagate** whenever you want to propagate the chosen settings to member chassis. For more information on propagation of leader chassis properties to member chassis, see the *Online Help*.

Server Inventory for MCM group

A group is a lead chassis that has 0 to 8 chassis group members. The **Chassis Group Health** page displays all the member chassis and allows you to save the server inventory report to a file, using standard browser download capability. The report contains data for:

- All servers currently in all the group chassis (including the leader).
- Empty slots and extension slots (including full height and double width server modules).

Saving Server Inventory Report

To save the server inventory report using the CMC web interface:

- 1. In the left pane, select the **Group**.
- 2. On the Chassis Group Health page, click Save Inventory Report. The File Download dialog box is displayed asking you to open or save the file.
- 3. Click Save and specify the path and file name for the server module inventory report.
 - NOTE: The chassis group leader and chassis group member chassis, and the server module in the associated chassis, must be turned on to get the most accurate server module inventory report.

Exported Data

The server inventory report contains data that was most recently returned by each Chassis Group member during the Chassis Group leader's normal polling (after every 30 seconds).

To get the most accurate server inventory report:

- The Chassis Group leader chassis and all Chassis Group member chassis must be in Chassis Power State On.
- All servers in the associated chassis must be turned on.

The inventory data for the associated chassis and servers may be missing in inventory report, if a subset of the Chassis Group member chassis are:

- In Chassis Power State Off state
- Turned off
- NOTE: If a server is inserted while the chassis is turned off, the model number is not displayed anywhere in the web interface until the chassis is turned on again.

The following table lists the specific data fields and specific requirements for fields to be reported for each server:

Table 16. Server Module Inventory Field Descriptions

Data Field	Example
Chassis Name	Data Center Chassis Leader
Chassis IP Address	192.168.0.1
Slot Location	1
Slot Name	SLOT-01
Host Name	Corporate Webserver

Table 16. Server Module Inventory Field Descriptions (continued)

Data Field	Example
	NOTE: Requires a Server Administrator agent running on the Server; otherwise shown as blank.
Operating System	Windows Server 2008 i NOTE: Requires a Server Administrator agent running on the Server; otherwise shown as blank.
Model	PowerEdgeM610
Service Tag	1PB8VF1
Total System Memory	4.0 GB i NOTE: Requires VRTX CMC 1.0 (or later) on member; otherwise shown as blank.
# of CPUs	2 (i) NOTE: Requires VRTX CMC 1.0 (or later) on member; otherwise shown as blank.
CPU Info	Intel (R) Xeon (R) CPU E5502 @1.87GHzn i NOTE: Requires VRTX CMC 1.0 (or later) on member; otherwise shown as blank.

Data Format

The inventory report is generated in a .CSV file format such that it can be imported to various tools, such as Microsoft Excel. The inventory report .CSV file can be imported into the template by selecting the **Data** > **From Text** in MS Excel. After the inventory report is imported into MS Excel, and if a message is displayed asking you for additional information, select comma-delimited to import the file into MS Excel.

Chassis Group Inventory and Firmware Version

The **Chassis Group Firmware Version** page displays the group inventory and firmware versions of the servers and the server components in the chassis. This page also enables you to organize the inventory information and filter the firmware versions view. The displayed view is based on the servers or any of the following chassis server components:

- BIOS
- iDRAC
- CPLD
- USC
- Diagnostics
- OS Drivers
- RAID
- NIC

NOTE: The inventory information displayed for the chassis group, member chassis, servers, and server components is updated every time a chassis is added or removed from the group.

Viewing Chassis Group Inventory

To view the chassis group using CMC web interface, in the left pane, select **Group**. Click **Properties** > **Firmware Version**. The **Chassis Group Firmware Version** page displays all the chassis in the group.

Viewing Selected Chassis Inventory Using Web Interface

To view the selected chassis inventory using CMC Web interface:

In the system tree, select Group. click Properties > Firmware Version.
 The Chassis Group Firmware Version page displays all the chassis in the group.

In the Select a Chassis section, select the member chassis for which you want to view the inventory.
 The Firmware View Filter section displays the server inventory for the selected chassis and the firmware versions of all the server components.

Viewing Selected Server Component Firmware Versions Using Web Interface

To view the firmware versions of selected server components using CMC web interface:

- In the left pane, select Group. Click Properties > Firmware Version.
 The Chassis Group Firmware Version page displays all the chassis in the group.
- 2. In the Select a Chassis section, select the member chassis for which you want to view the inventory.
- 3. In the Firmware View Filter section, select Components.
- 4. In the Components list, select the required component- BIOS, iDRAC, CPLD, USC, Diagnostics, OS Drive, RAID devices (up to 2), and NIC devices (up to 6), for which you want to view the firmware version.
 The firmware versions of the selected component for all the servers in the selected member chassis are displayed.

Chassis Configuration Profiles

The Chassis Configuration Profiles feature enables you to configure the chassis with the chassis configuration profiles stored in the network share or local management station, and also restore configuration of the chassis.

To access the **Chassis Configuration Profiles** page in the CMC web interface, in the system tree, go to **Chassis Overview** and click **Setup** > **Profiles**. The **Chassis Configuration Profiles** page is displayed.

You can perform the following tasks by using the Chassis Configuration Profiles feature:

- Configure a chassis using chassis configuration profiles in local management station for initial configuration.
- Save the current chassis configuration settings to an XML file on the network share or local management station.
- Restore the chassis configuration.
- Import chassis profiles (XML files) to the network share from a local management station.
- Export chassis profiles (XML files) from the network share to a local management station.
- Apply, edit, delete, or export a copy of the profiles stored on the network share.

Saving Chassis Configuration

You can save the current chassis configuration to an XML file on a network share or local management station. The configurations include all the chassis properties that can be modified using the CMC web interface and RACADM commands. You can also use the XML file that is saved to restore the configuration on the same chassis or to configure other chassis.

(i) NOTE: Server and iDRAC settings are not saved or restored with the chassis configuration.

To save the current chassis configuration, perform the following tasks:

- Go to the Chassis Configuration Profiles page. In the Save and Backup > Save Current Configuration section, enter a name for the profile in the Profile Name field.
 - NOTE: While saving the current chassis configuration, the standard ASCII extended character set is supported. However, the following special characters are not supported:

", .,
$$*$$
, >, <, \, /, :, and |

- 2. Select one of the following profile types from the **Profile Type** option:
 - **Replace** Includes attributes of the entire CMC configuration except write-only attributes such as user passwords and service tags. This profile type is used as a backup configuration file to restore the complete chassis configuration including identity information such as IP addresses.
 - Clone Includes all the **Replace** type profile attributes. The Identity attributes such as MAC address and IP address are commented out for safety reasons. This profile type is used to clone a new chassis.

- 3. Select one of the following locations from the **Profile Location** drop-down menu to store the profile:
 - Local To save the profile in the local management station.
 - **Network Share** To save the profile in a shared location.
- 4. Click **Save** to save the profile to the selected location.

After the action is complete, the Operation Successful message is displayed:

NOTE: To view the settings that are saved to the XML file, in the **Stored Profiles** section, select the saved profile and click **View** in the **View Profiles** column.

Restoring Chassis Configuration Profile

You can restore the configuration of a chassis by importing the backup file (.xml or .bak) on the local management station or the network share to which the chassis configurations were saved. The configurations include all the properties available through the CMC web interface, RACADM commands, and settings.

To restore the chassis configuration, perform the following tasks:

- 1. Go to the Chassis Configuration Profiles page. In the Restore Configuration > Restore Chassis Configuration section, click Browse and select the backup file to import the saved chassis configuration.
- 2. Click **Restore Configuration** to upload an encrypted backup file (.bak) or a .xml stored profile file to the CMC. The CMC web interface returns to the login page after a successful restore operation.
- NOTE: If the backup files (.bak) of the earlier versions of CMC, are loaded on the latest version of CMC where FIPS is enabled, reconfigure all the 16 CMC local user passwords. However, the password of the first user is reset to "calvin".
- NOTE: When a chassis configuration profile is imported from a CMC, which does not support the FIPS feature, to a CMC where FIPS is enabled, the FIPS remains enabled in the CMC.
- NOTE: If you change the FIPS mode in the chassis configuration profile, the DefaultCredentialMitigation is enabled.

Viewing Stored Chassis Configuration Profiles

To view the chassis configuration profiles stored on the network share, go to the **Chassis Configuration Profiles** page. In the **Chassis Configuration Profiles** > **Stored Profiles** section, select the profile and click **View** in the **View Profile** column. The **View Settings** page is displayed. For more information on the displayed settings, see the *CMC Online Help*.

Applying Chassis Configuration Profiles

You can apply chassis configuration to the chassis if the chassis configuration profiles are available as stored profiles on the network share. To initiate a chassis configuration operation, you can apply a stored profile to a chassis.

To apply a profile to a chassis, perform the following tasks:

- 1. Go to the **Chassis Configuration Profiles** page. In the **Stored Profiles** section, select the stored profile that you want to apply.
- 2. Click Apply Profile.

A warning message is displayed that applying a new profile overwrites the current settings and also reboots the selected chassis. You are prompted to confirm if you want to continue the operation.

3. Click **OK** to apply the profile to the chassis.

Exporting Chassis Configuration Profiles

You can export chassis configuration profiles that are saved on the network share to a specified path on a management station.

To export a stored profile, perform the following tasks:

1. Go to the Chassis Configuration Profiles page. In the Chassis Configuration Profiles > Stored Profiles section, select the required profile and then click Export Copy of Profile.

A File Download message is displayed prompting you to open or save the file.

2. Click Save or Open to export the profile to the required location.

Editing Chassis Configuration Profiles

You can edit chassis configuration profile name of a chassis.

To edit a chassis configuration profile name, perform the following tasks:

- Go to the Chassis Configuration Profiles page. In the Chassis Configuration Profiles > Stored Profiles section, select
 the required profile and then click Edit Profile.
 - The **Edit Profile** window is displayed.
- Enter a desired profile name in the Profile Name field and click Edit Profile.
 Operation Successful message is displayed.
- 3. Click OK.

Deleting Chassis Configuration Profiles

You can delete a chassis configuration profile that is stored on the network share.

To delete a chassis configuration profile, perform the following tasks:

- Go to the Chassis Configuration Profiles page. In the Chassis Configuration Profiles > Stored Profiles section, select
 the required profile and then click Delete Profile.
 - A warning message is displayed indicating that deleting a profile would delete the selected profile permanently.
- 2. Click **OK** to delete the selected profile.

Configuring Multiple CMCs Using RACADM

Using RACADM, you can configure one or more CMCs with identical properties.

When you query a specific CMC card using its group ID and object ID, RACADM creates the racadm.cfg configuration file from the retrieved information. By exporting the file to one or more CMCs, you can configure your controllers with identical properties in a minimal amount of time.

- NOTE: Some configuration files contain unique CMC information (such as the static IP address) that must be modified before you export the file to other CMCs.
- 1. Use RACADM to query the target CMC that contains the desired configuration.
 - NOTE: The generated configuration file is myfile.cfg. You can rename the file. The .cfg file does not contain user passwords. When the .cfg file is uploaded to the new CMC, you must re-add all passwords.
- 2. At the command prompt, type:

```
racadm getconfig -f myfile.cfg
```

- NOTE: Redirecting the CMC configuration to a file using getconfig -f is only supported with the remote RACADM interface.
- 3. Modify the configuration file using a plain-text editor (optional). Any special formatting characters in the configuration file may corrupt the RACADM database.
- 4. Use the newly created configuration file to modify a target CMC. At the command prompt, type:

```
racadm config -f myfile.cfg
```

5. Reset the target CMC that was configured. At the command prompt, type:

```
racadm reset
```

The getconfig -f myfile.cfg subcommand requests the CMC configuration for the active CMC and generates the myfile.cfg file. If required, you can rename the file or save it to a different location.

You can run the getconfig command to perform the following actions:

- Display all configuration properties in a group (specified by group name and index).
- Display all configuration properties for a user by user name.

The config subcommand loads the information into other CMCs. The Server Administrator uses the config command to synchronize the user and password database.

Creating a CMC Configuration File

The CMC configuration file, <filename>.cfg, is used with the racadm config -f <filename>.cfg command to create a simple text file. The command allows you to build a configuration file (similar to a .ini file) and configure the CMC from this file.

You may use any file name, and the file does not require a .cfg extension (although it is referred to by that designation in this subsection).

NOTE: For more information about the getconfig subcommand, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

RACADM parses the .cfg file when it is first loaded on to the CMC to verify that a valid group and object names are present, and that simple syntax rules are being followed. Errors are flagged with the line number that detected the error, and a message explains the problem. The entire file is parsed for correctness, and all errors display. If an error is found in the .cfg file, write commands are not transmitted to the CMC. You must correct all errors before any configuration can take place.

To check for errors before you create the configuration file, use the -c option with the config subcommand. With the -c option, config only verifies syntax and does not write to the CMC.

Follow these guidelines when you create a .cfg file:

- If the parser encounters an indexed group, it is the value of the anchored object that differentiates the various indexes.
 - The parser reads in all of the indexes from the CMC for that group. Any objects within that group are modifications when the CMC is configured. If a modified object represents a new index, the index is created on the CMC during configuration.
- You cannot specify a desired index in a .cfg file.
 - Indexes may be created and deleted. Over time, the group may become fragmented with used and unused indexes. If an index is present, it is modified. If an index is not present, the first available index is used.
 - This method allows flexibility when adding indexed entries where you do not need to make exact index matches between all the CMCs being managed. New users are added to the first available index. A .cfg file that parses and runs correctly on one CMC may not run correctly on another, if all indexes are full and you must add a new user.
- Use the racresetcfg subcommand to configure both the CMCs with identical properties.
 - Use the racresetcfg subcommand to reset the CMC to original defaults, and then run the racadm config f <filename>.cfg command. Make sure that the .cfg file includes all desired objects, users, indexes, and other parameters. For a complete list of objects and groups, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.
 - CAUTION: Use the racresetcfg subcommand to reset the database and the CMC Network Interface settings to the original default settings and remove all users and user configurations. While the root user is available, other users' settings are also reset to the default settings.
- If you type racadm getconfig -f <filename> .cfg, the command builds a .cfg file for the current CMC configuration. This configuration file can be used as an example and as a starting point for your unique .cfg file.

Parsing Rules

- Lines that start with a hash character (#) are treated as comments.
 - A comment line must start in column one. A "#" character in any other column is treated as a # character.
 - Some modem parameters may include # characters in their strings. An escape character is not required. You may want to generate a .cfg from a racadm getconfig -f <filename> .cfg command, and then perform a racadm config -f <filename> .cfg command to a different CMC, without adding escape characters.

For example:

```
#
# This is a comment
[cfgUserAdmin]
cfgUserAdminPageModemInitString= <Modem init # not
a comment>
```

All group entries must be surrounded by open- and close-brackets ([and]).

The starting [character that denotes a group name must be in column one. This group name must be specified before any of the objects in that group. Objects that do not include an associated group name generate an error. The configuration data is organized into groups as defined in the database property chapter of the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide*. The following example displays a group name, object, and the object's property value:

```
[cfgLanNetworking] -{group name}
cfgNicIpAddress=143.154.133.121 {object name}
{object value}
```

• All parameters are specified as "object=value" pairs with no white space between the object, =, or value. White spaces that are included after the value are ignored. A white space inside a value string remains unmodified. Any character to the right of the = (for example, a second =, a #, [,], and so on) is taken as-is. These characters are valid modem chat script characters.

```
[cfgLanNetworking] -{group name}
cfgNicIpAddress=143.154.133.121 {object value}
```

The .cfg parser ignores an index object entry.

You cannot specify which index is used. If the index already exists, it is either used or the new entry is created in the first available index for that group.

The racadm getconfig -f <filename>.cfg command places a comment in front of index objects, allowing you to see the included comments.

NOTE: You may create an indexed group manually using the following command:

```
\label{local_confidence} \mbox{racadm config -g <group name>} \mbox{ -o <anchored object> -i <index 1-4> <unique anchored name>}
```

• The line for an indexed group cannot be deleted from a .cfg file. If you do delete the line with a text editor, RACADM stops when it parses the configuration file and alert you of the error.

You must remove an indexed object manually using the following command:

```
racadm config -g <groupname> -o <objectname> -i <index 1-4> ""
```

i NOTE: A NULL string (identified by two " characters) directs the CMC to delete the index for the specified group.

To view the contents of an indexed group, run the following command:

```
racadm getconfig -g <groupname> -i <index 1-4>
```

• For indexed groups the object anchor must be the first object after the [] pair. The following are examples of the current indexed groups:

```
[cfgUserAdmin]
cfgUserAdminUserName= <USER_NAME>
```

• When using remote RACADM to capture the configuration groups into a file, if a key property within a group is not set, the configuration group is not saved as part of the configuration file. If these configuration groups are needed to be cloned onto other CMCs, the key property must be set before executing the getconfig -f command. Alternatively, you can manually enter the missing properties into the configuration file after running the getconfig -fcommand. This is true for all the RACADM-indexed groups.

This is the list of the indexed groups that exhibit this behavior and their corresponding key properties:

o cfgUserAdmin — cfgUserAdminUserName

- $\circ \quad \mathsf{cfgEmailAlert} \mathsf{cfgEmailAlertAddress}$
- o cfgTraps cfgTrapsAlertDestlPAddr
- o cfgStandardSchema cfgSSADRoleGroupName
- o cfgServerInfo cfgServerBmcMacAddress

Modifying the CMC IP Address

When you modify the CMC IP address in the configuration file, remove all unnecessary <variable> = <value> entries. Only the actual variable group's label with [and] remains, including the two <variable> = <value> entries pertaining to the IP address change.

Example:

```
#
# Object Group "cfgLanNetworking"
#
[cfgLanNetworking]
cfgNicIpAddress=10.35.10.110
cfgNicGateway=10.35.10.1
```

This file is updated as follows:

```
#
# Object Group "cfgLanNetworking"
#
[cfgLanNetworking]
cfgNicIpAddress=10.35.9.143
# comment, the rest of this line is ignored
cfgNicGateway=10.35.9.1
```

The command racadm config -f <myfile>.cfg parses the file and identifies any errors by line number. A correct file updates the proper entries. Additionally, you can use the same getconfig command from the previous example to confirm the update.

Use this file to download company-wide changes or to configure new systems over the network with the command, racadm getconfig -f <myfile>.cfg.

(i) NOTE: Anchor is a reserved word and should not be used in the .cfg file.

Configuring Multiple CMCs through RACADM Using Chassis Configuration Profiles

By using chassis configuration profiles, you can export the chassis configuration profiles as an XML file and import it to another chassis.

Use RACADM get command for export operation and set command for import operation. You can export chassis profiles (XML files) from CMC to the network share or to a local management station and import chassis profiles (XML files) from the network share or from a local management station.

i NOTE: By default, the export is done as clone type. You can use the --clone to get the clone type profile in XML file.

The import and export operation to and from the network share can be done through local RACADM as well as remote RACADM. Whereas, the import and export operation to and from the local management can be done only through remote RACADM interface.

Exporting Chassis Configuration profiles

You can export chassis configuration profiles to network share by using the get command.

 To export the chassis configuration profiles as clone.xml file to CIFS network share by using get command, type the following:

```
racadm get -f clone.xml -t xml -l //xx.xx.xx.xx/PATH -u USERNAME -p PASSWORDCMC
```

2. To export the chassis configuration profiles as clone.xml file to NFS network share by using get command, type the following:

```
racadm get -f clone.xml -t xml -l xx.xx.xx:/PATH
```

You can export chassis configuration profiles to network share through a remote RACADM interface.

1. To export the chassis configuration profiles as clone.xml file to CIFS network share, type the following:

```
racadm -r xx.xx.xx -u USERNAMECMC -p PASSWORDCMC get -f clone.xml -t xml -l // xx.xx.xx.xx/PATH -u USERNAME -p PASSWORD
```

2. To export the chassis configuration profiles as clone.xml file to NFS network share, type the following:

```
racadm -r xx.xx.xx.xx -u USERNAMECMC -p PASSWORDCMC get -f clone.xml -t xml -l xx.xx.xx.xx:/PATH
```

you can export chassis configuration profiles to local management station through remote RACADM interface.

1. To export the chassis configuration profiles as clone.xml file, type the following:

```
racadm -r xx.xx.xx.xx -u USERNAMECMC -p PASSWORDCMC get -f clone.xml -t xml
```

Importing Chassis Configuration profiles

You can import chassis configuration profiles from network share to another chassis by using the set command.

1. To import the chassis configuration profiles from CIFS network share, type the following:

```
racadm set -f clone.xml -t xml -l //xx.xx.xx.xx/PATH -u USERNAME -p PASSWORDCMC
```

2. To import the chassis configuration profiles from NFS network share, type the following:

```
racadm set -f clone.xml -t xml -l xx.xx.xx:/PATH
```

You can import chassis configuration profiles from network share through remote RACADM interface.

1. To import the chassis configuration profiles from CIFS network share, type the following:

```
racadm -r xx.xx.xx.xx -u USERNAMECMC -p PASSWORDCMC set -f clone.xml -t xml -l // xx.xx.xx.xx/PATH -u USERNAME -p PASSWORD
```

2. To import the chassis configuration profiles from NFS network share, type the following:

```
racadm -r xx.xx.xx -u USERNAMECMC -p PASSWORDCMC set -f clone.xml -t xml -l xx.xx.xx.xx:/PATH
```

You can import chassis configuration profiles from local management station through remote RACADM interface.

1. To export the chassis configuration profiles as clone.xml file, type the following:

```
racadm -r xx.xx.xx -u USERNAMECMC -p PASSWORDCMC set -f clone.xml -t xml
```

Parsing Rules

You can manually edit properties of an exported XML file of chassis configuration profiles.

An XML file contains the following properties:

- System Configuration, which is the parent node.
- component, which is the primary child node.

 Attributes, which contains name and value. You can edit these fields. For example, you can edit the Asset Tag value as follows:

```
<attribute Name="ChassisInfo.1#AssetTag">xxxxxx</attribute>
```

Example of an XML file is as follows:

Viewing and Ending CMC Sessions

You can view the number of users currently logged in to iDRAC7 and end the user sessions.

i NOTE: To end a session, you must have Chassis Configuration Administrator privilege.

Viewing and Ending CMC Sessions Using Web Interface

To view or end a session using the web interface:

1. In the left pane, go to Chassis Overview and click Network > Sessions.

The **Sessions** page displays the session ID, username, IP address, and session type. For more information about these properties, see the *Online Help*.

2. To end the session, click **Terminate** for a session.

Viewing and Ending CMC Sessions Using RACADM

You must have administrator privileges to end CMC sessions using RACADM.

To view the current user sessions, use the getssninfo command.

To end a user session, use the closessn command.

For more information about these commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/support/manuals**.

Configuring Servers

You can configure the following settings of a server:

- Slot Names
- iDRAC Network Settings
- DRAC Virtual LAN Tag Settings
- First Boot Device
- Server FlexAddress
- Remote File Share
- BIOS Settings Using Server Clone

Topics:

- Configuring Slot Names
- Configuring iDRAC Network Settings
- Configuring iDRAC Virtual LAN Tag Settings
- Setting First Boot Device
- Configuring Server FlexAddress
- Configuring Remote File Share
- Configuring Profile Settings Using Server Configuration Replication

Configuring Slot Names

Slot names are used to identify individual servers. When choosing slot names, the following rules apply:

- Names may contain a maximum of 15 non-extended ASCII characters (ASCII codes 32 through 126). Standard and special
 characters are allowed in the names.
- Slot names must be unique within the chassis. Slots should not have the same name.
- Strings are not case-sensitive. Server-1, server-1, and SERVER-1 are equivalent names.
- Slot names must not begin with the following strings:
 - o Switch-
 - o Fan-
 - o PS-
 - o DRAC-
 - o MC-
 - o Chassis
 - o Housing-Left
 - o Housing-Right
 - o Housing-Center
- The strings Server-1 through Server-4 may be used, but only for the corresponding slot. For example, Server-3 is a valid name for slot 3, but not for slot 4. However, Server-03 is a valid name for any slot.
 - i NOTE: To change a slot name, you must have the Chassis Configuration Administrator privilege.

The slot name setting in the web interface resides on CMC only. If a server is removed from the chassis, the slot name setting does not remain with the server.

The slot name setting in the CMC web interface always overrides any change you make to the display name in the iDRAC interface.

To edit a slot name using the CMC Web interface:

- 1. In the left pane, go to Chassis Overview > Server Overview > Setup > Slot Names.
- 2. On the Slot Names page, edit the slot name, in the Slot Name field.

- 3. To use a server's host name as slot name, select the **Use Host Name for the Slot name** option. This overrides the static slot names with the server's Host Name (or system name), if available. This requires the OMSA agent to be installed on the server. For more information about the OMSA agent, see the *Dell OpenManage Server Administrator User's Guide* available at dell.com/support/manuals.
- **4.** To use the iDRAC DNS name as slot name, select the **Use iDRAC DNS Name for Slot Name** option. This option replaces the static slot names with the respective iDRAC DNS names, if it is available. If iDRAC DNS names are not available, the default or edited slot names are displayed.
 - NOTE: To use the Use iDRAC DNS Name for Slot Name option, you must have the Chassis Configuration Administrator privilege.
- 5. To save the settings, click Apply.

To restore the default slot name (SLOT-01 through SLOT-04) on the basis of a server's slot position) to a server, click **Restore Default Value**.

Configuring iDRAC Network Settings

To use this feature, you must have an Enterprise License. You can configure the iDRAC network configuration setting of a server. You can use the QuickDeploy settings to configure the default iDRAC network configuration settings and root password for severs that are installed later. These default settings are the iDRAC QuickDeploy settings.

For more information about iDRAC, see the iDRAC User's Guide at dell.com/support/manuals.

Configuring iDRAC QuickDeploy Network Settings

Use the QuickDeploy Settings to configure the network settings for newly inserted servers.

To enable and set the iDRAC QuickDeploy settings:

- 1. In the left pane, click Server Overview > Setup > iDRAC.
- 2. On the **Deploy iDRAC** page, in the **QuickDeploy Settings** section, specify the settings mentioned in the following table. For more information about the fields, see the *Online Help*.

Table 17. QuickDeploy Settings

Setting	Description
Action When Server is Inserted	 Select one of the following options from the list: No Action — The action is not performed when the server is inserted. QuickDeploy Only — Select this option to apply iDRAC network settings when a new server is inserted in the chassis. The specified autodeployment settings are used to configure the new iDRAC, which includes the root user password if Change Root Password is selected. Server Profile Only — Select this option to apply server profile assigned when a new server is inserted in the chassis. QuickDeploy and Server Profile — Select this option to first apply the iDRAC network settings, and then to apply the server profile assigned when a new server is inserted in the chassis.
Set iDRAC Root Password on Server Insertion	Select the option to change iDRAC root password to match the value provided in the iDRAC Root Password field, when a server is inserted.
iDRAC Root Password	When the Set iDRAC Root Password on Server Insertion and QuickDeploy Enabled options are selected, this password value is assigned to a server's iDRAC root user password when the server is inserted into a chassis. The password can have 1 to 20 printable (including white spaces) characters.
Confirm iDRAC Root Password	Allows you to retype the password provided in the Password field.
Enable iDRAC LAN	Enables or disables the iDRAC LAN channel. By default, this option is cleared.
Enable iDRAC IPv4	Enables or disables IPv4 on iDRAC. By default, this option is selected.

Table 17. QuickDeploy Settings (continued)

Setting	Description
Enable iDRAC IPMI over LAN	Enables or disables the IPMI over LAN channel for each iDRAC present in the chassis. By default, this option is selected.
Enable iDRAC IPv4 DHCP	Enables or disables DHCP for each iDRAC present in the chassis. If this option is enabled, the fields QuickDeploy IP , QuickDeploy Subnet Mask , and QuickDeploy Gateway are disabled, and cannot be modified since DHCP is used to automatically assign these settings for each iDRAC. To select this option, you must select the Enable iDRAC IPv4 option. Quick Deploy IP address is provided with two options — 2 and 4.
Starting iDRAC IPv4 Address (Slot 1)	Specifies the static IP address of iDRAC in the server, in slot 1 of the enclosure. The IP address of each subsequent iDRAC is incremented by 1 for each slot from slot 1's static IP address. In the case where the IP address plus the slot number is greater than the subnet mask, an error message is displayed. NOTE: The subnet mask and the gateway are not incremented such as the IP address.
	For example, if the starting IP address is 192.168.0.250 and the subnet mask is 255.255.0.0 then the QuickDeploy IP address for slot 15 is 192.168.0.265. If the subnet mask is 255.255.255.0, the QuickDeploy IP address range is not fully within QuickDeploy Subnet error message is displayed when you click Save QuickDeploy Settings or Auto-Populate Using QuickDeploy Settings.
iDRAC IPv4 Netmask	Specifies the QuickDeploy subnet mask that is assigned to all newly inserted servers.
iDRAC IPv4 Gateway	Specifies the QuickDeploy default gateway that is assigned to all the DRAC present in the chassis.
Enable iDRAC IPv6	Enables IPv6 addressing for each iDRAC present in the chassis that is IPv6 capable.
Enable iDRAC IPv6 Autoconfiguration	Enables the iDRAC to obtain IPv6 settings (address and prefix length) from a DHCPv6 server and also enables stateless address auto configuration. By default, this option is enabled.
iDRAC IPv6 Gateway	Specifies the default IPv6 gateway to be assigned to the iDRACs. The default value is "::".
iDRAC IPv6 Prefix Length	Specifies the prefix length to be assigned for the IPv6 addresses on the iDRAC. The default value is 64.
Use CMC DNS Settings	Communicates the CMC DNS server settings (IPv4 and IPv6) to iDRAC when a blade server is inserted in the chassis.

3. Click Save QuickDeploy Settings to save the settings. If you have made changes to the iDRAC network setting, click Apply iDRAC Network Settings to deploy the settings to the iDRAC.

The QuickDeploy feature only executes when it is enabled, and a server is inserted in the chassis. If **Set iDRAC Root Password on Server Insertion** and **QuickDeploy Enabled** are enabled, the user is prompted using the LCD interface to allow or not allow the password change. If there are network configuration settings that differ from the current iDRAC settings, the user is prompted to either accept or reject the changes.

NOTE: If there is a LAN or IPMI over LAN difference, the user is prompted to accept the QuickDeploy IP address setting. If the difference is the DHCP setting, the user is prompted to accept the DHCP QuickDeploy setting.

To copy the QuickDeploy settings into the iDRAC Network Settings section, click Auto-Populate Using QuickDeploy Settings. The QuickDeploy network configurations settings are copied into the corresponding fields in the iDRAC Network Configuration Settings table.

NOTE: Changes made to QuickDeploy fields are immediate, but changes made to one or more iDRAC server network configuration settings may require a couple of minutes to propagate from CMC to iDRAC. Clicking **Refresh** early may display only partially correct data for one or more iDRAC servers.

Assigning QuickDeploy IP Address to Servers

The figure here shows the QuickDeploy IP addresses assignment to the servers when there are four half-height servers in VRTX chassis:



The following figure shows the QuickDeploy IP addresses assignment to the servers when there are two full-height blades in VRTX chassis:



Modifying iDRAC Network Settings for Individual Server iDRAC

Using this feature, you can configure the iDRAC network configurations settings for each installed server. The initial values displayed for each of the fields are the current values read from the iDRAC. To use this feature, you must have an Enterprise License.

To modify the iDRAC Network Settings:

- 1. In the left pane, click **Server Overview**, and then click **Setup**. On the **Deploy iDRAC** page the **iDRAC Network Settings** section lists the iDRAC IPv4 and IPv6 network configuration settings of all the installed servers.
- 2. Modify the iDRAC network settings as required for the server(s).
 - NOTE: You must select the **Enable LAN** option to specify the IPv4 or IPv6 settings. For information about the fields, see the *Online Help*.
- 3. To deploy the setting to iDRAC, click **Apply iDRAC Network Settings**. Any changes made to the **QuickDeploy Settings** are also saved.

The **iDRAC Network Settings** table reflects future network configuration settings; the values shown for installed servers may or may not be the same as the currently installed iDRAC network configuration settings. Click **Refresh** to update the **iDRAC Deploy** page with each installed iDRAC network configuration settings after changes are made.

NOTE: Changes made to QuickDeploy fields are immediate, but changes made to one or more iDRAC server network configuration settings may require a couple of minutes to propagate from CMC to iDRAC. Clicking **Refresh** too soon may display only partially correct data for a one or more iDRAC servers.

Modifying iDRAC Network Settings Using RACADM

RACADM config or getconfig commands support the -m <module> option for the following configuration groups:

- cfgLanNetworking
- cfgIPv6LanNetworking
- cfgRacTuning
- cfgRemoteHosts
- cfgSerial
- cfgSessionManagement

For more information about the property default values and ranges, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/support/manuals**.

Configuring iDRAC Virtual LAN Tag Settings

Virtual LAN (VLAN) tags enable multiple VLANs to coexist on the same physical network cable and to segregate the network traffic for security or load management purposes. When you enable the VLAN functionality, each network packet is assigned a VLAN tag. VLAN tags are chassis properties. They remain with the chassis even when a component is removed.

NOTE: The VLAN ID configured using CMC is applied to iDRAC only when iDRAC is in dedicated mode. If iDRAC is in shared LOM mode, the VLAN ID changes made in iDRAC are not displayed on the CMC GUI.

Configuring iDRAC Virtual LAN Tag Settings Using RACADM

• Specify the Virtual LAN ID and priority of a particular server with the following command:

```
racadm setniccfg -m server-<n> -v <VLAN id> <VLAN priority>
```

The valid values for $\langle n \rangle$ are 1–4.

The valid values for <VLAN> are 1-4000 and 4021-4094. Default is 1.

The valid values for $\langle VLAN \ priority \rangle$ are 0-7. Default is 0.

For example:

```
racadm setniccfg -m server-1 -v 1 7
```

For example:

• To remove a server VLAN, disable the VLAN capabilities of the specified server's network:

```
racadm setniccfg -m server-<n> -v
```

The valid values for $\langle n \rangle$ are 1–4.

For example:

racadm setniccfg -m server-1 -v

Configuring iDRAC Virtual LAN Tag Settings Using Web Interface

To configure Vitual LAN(VLAN) for server:

- 1. Go to any of the following pages:
 - In the left pane, click Chassis Overview > Network > VLAN.
 - In the left pane, click Chassis Overview > Server Overview and click Setup > VLAN.
- 2. On the **VLAN Tag Settings** page, in the **iDRAC** section, enable VLAN for the server(s), set the priority and enter the ID. For more information about the fields, see the *Online Help*.
- **3.** Click **Apply** to save the settings.

Setting First Boot Device

You can specify the CMC first boot device for each server. This may not be the actual first boot device for the server, or may not even represent a device present in that server. It represents a device sent by CMC to the server and used as its first boot

device of that server. This device can be set as the default first-boot device or an one-time device so that you can boot an image to perform tasks such as running diagnostics or reinstalling an operating system.

You can set the first boot device for the next boot only or for all subsequent reboots. You can also set the first boot device for the server. The system boots from the selected device on the next and subsequent reboots and remains as the first boot device in the BIOS boot order, until it is changed again either from the CMC web interface (**Chassis Overview** > **Server Overview** > **Setup** > **First Boot Device**) or from the BIOS boot sequence.

i NOTE: The first boot device setting in CMC web Interface overrides the system BIOS boot settings.

The boot device that you specify must exist and contain a bootable media.

You can set the following devices for first boot.

Table 18. Boot Devices

Boot Device	Description
PXE	Boot from a Preboot Execution Environment (PXE) protocol on the Network Interface Card.
Hard Drive	Boot from the hard disk drive on the server.
Local CD/DVD	Boot from a CD or DVD drive on the server.
Virtual Floppy	Boot from the virtual floppy drive. The floppy drive (or a floppy disk image) is on another computer on the management network, and is attached using the iDRAC GUI console viewer.
Virtual CD/DVD	Boot from a virtual CD or DVD drive or CD or DVD ISO image. The optical drive or ISO image file is located on another computer or boot disk available on the management network and is attached using the iDRAC GUI console viewer.
Local SD Card	Boot from the local SD (Secure Digital) card—for servers that support iDRAC 6 and iDRAC 7 systems only.
Local Floppy	Boot from a floppy disk in the local floppy disk drive.
Remote File Share	Boot from a Remote File Share (RFS) image. The image file is attached using the iDRAC GUI console viewer.

Setting First Boot Device For Multiple Servers Using CMC Web Interface

NOTE: To set the first boot device for servers, you must have the Server Administrator privileges or Chassis Configuration Administrator privileges, and the iDRAC login privileges.

To set the first boot device for multiple servers:

- 1. In the left pane, click Server Overview > Setup > First Boot Device. A list of servers is displayed.
- 2. In the **First Boot Device** column, from the drop-down menu corresponding to a server, select the boot device you want to use for a server.
- 3. If you want the server to boot from the selected device every time it boots, clear the **Boot Once** option for the server. If you want the server to boot from the selected device only on the next boot cycle, select the **Boot Once** option for the server.
- 4. Click Apply to save the settings.

Setting First Boot Device For Individual Server Using CMC Web Interface

NOTE: To set the first boot device for servers, you must have Server Administrator privileges or Chassis Configuration Administrator privileges and iDRAC login privileges.

To set the first boot device for individual servers:

- 1. In the left pane, click Server Overview, and then click the server for which you want to set the first boot device.
- 2. Go to Setup > First Boot Device. The First Boot Device page is displayed.

- 3. From the First Boot Device drop-down menu, select the boot device you want to use for each server.
- 4. If you want the server to boot from the selected device every time it boots, clear the Boot Once option for the server.
 If you want the server to boot from the selected device only on the next boot cycle, select the Boot Once option for the server
- 5. Click **Apply** to save the settings.

Setting First Boot Device Using RACADM

To set the first boot device, use the cfgServerFirstBootDevice object.

To enable boot once for a device, use the cfgServerBootOnce object.

For more information about these objects, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Configuring Server FlexAddress

For information about configuring FlexAddress for servers, see Configuring FlexAddress for Chassis-Level Fabric and Slots Using CMC Web Interface. To use this feature, you must have an Enterprise License.

Configuring Remote File Share

The Remote Virtual Media File Share feature maps a file from a share drive on the network to one or more servers through CMC to deploy or update an operating system. When connected, the remote file is accessible similar to a file that you can access on a local server. Two types of media are supported: floppy drives and CD/DVD drives.

To perform a remote file share operation (connect, disconnect, or deploy), you must have the **Chassis Configuration Administrator** or **Server Administrator** privileges. To use this feature, you must have an Enterprise license.

To configure the remote file share:

- 1. In the left pane, click Server Overview > Setup > Remote File Share.
- 2. On the **Deploy Remote File Share** page, type appropriate data in the fields. For more information about the field descriptions, see the *Online Help*.
- **3.** To connect to a remote file share, click **Connect**. To connect a remote file share, you must provide the path, user name, and password. A successful operation allows access to the media.

Click **Disconnect** to disconnect a previously-connected remote file share.

Click **Deploy** to deploy the media device.

NOTE: Before you click the **Deploy** button, make sure that you save all the working files, because this action restarts the server.

When you click **Deploy**, the following tasks are executed:

- The remote file share is connected.
- The file is selected as the first boot device for the servers.
- The server is restarted.
- Power is supplied to the server if the server is turned off.

Configuring Profile Settings Using Server Configuration Replication

The server configurations replicating feature allows you to apply all profile settings from a specified server to one or more servers. Profile settings that can be replicated are those profile settings which can be modified and are intended to be replicated across servers. The following three profile groups for servers are displayed and can be replicated:

- BIOS This group includes only the BIOS settings of a server. These profiles are generated from CMC for PowerEdge VRTX version 1.00 and later.
- BIOS and Boot This group includes the BIOS and the Boot settings of a server. These profiles are generated from CMC for PowerEdge VRTX version 1.00 and later.
- All Settings This version includes all the settings of the server and components on that server. These profiles are generated from
 - o CMC for PowerEdge VRTX version 1.00 and later
 - o 12th generation servers with iDRAC7 1.00.00 or later and Lifecycle Controller 2 version 1.1 or later
 - o 13th generation servers with iDRAC8 with Lifecycle Controller 2.00.00.00 or later

The server configurations replication feature supports iDRAC7 and later servers. Earlier generation RAC servers are listed but are grayed out on the main page, and are not enabled to use this feature.

To use the server configurations replication feature:

- iDRAC must have the minimum version that is required.
- Server must be powered on.

You can:

- View profile settings on a server or from a saved profile.
- Save a profile from a server.
- Apply a profile to other servers.
- Import stored profiles from a management station or remote file share.
- Edit the profile name and description.
- Export stored profiles to a management station or remote file share.
- Delete stored profiles.
- Deploy selected profiles to the target devices using **Quick Deploy** option.
- Display the log activity for recent server profile tasks.

Accessing Server Profiles Page

You can add, manage, and apply server profiles to one or more servers using the Server Profiles page.

To access the **Server Profiles** page using the CMC web interface, in the left pane, go to **Chassis Overview** > **Server Overview**. Click **Setup** > **Profiles**. The **Server Profiles** page is displayed.

Adding or Saving Profile

Before copying the properties of a server, first capture the properties to a stored profile. Create a stored profile and provide a name and optional description for each profile. You can save a maximum of 16 Stored Profiles on the CMC nonvolatile extended storage media.

NOTE: If a remote share is available, you can store a maximum of 100 profiles using the CMC extended storage and remote share. For more information about the remote share, see Configuring Network Share Using CMC Web Interface.

Removing or disabling the nonvolatile extended storage media prevents access to stored profile and disables the Server Configuration Replication feature.

To add or save a profile:

- 1. Open the Server Profiles page. In the Server Profiles section, click Apply and Save Profiles.
- 2. Select the server from whose settings you want to generate the profile, and then click **Save Profile**. The **Save Profile** section is displayed.
- 3. Select Extended Storage or Network Share as the location to save the profile.
 - NOTE: The Network Share option is enabled and the details are displayed in the Stored Profiles section only if the network share is mounted and is accessible. If the Network Share is not connected, configure the Network Share for the chassis. To configure the Network Share, click Edit in the Stored Profiles section. For more information about configuring the network share, see Configuring Network Share Using CMC Web Interface.
- 4. In the Profile Name and Description fields, type the profile name and description (optional), and then click Save Profile.
 - NOTE: When saving a Server Profile, the standard ASCII extended character set is supported. However, the following special characters are not supported:

CMC communicates with the LC to get the available server profile settings and store them as a named profile.

A progress indicator indicates that the Save operation is in progress. After the action is complete, a message, "Operation Successful" is displayed.

NOTE: The process to gather the settings runs in the background. Hence, it may take some time before the new profile is displayed. If the new profile is not displayed, check the profile log for errors.

Applying Profile

Server configuration replication is possible only when server profiles are available as stored profiles in the nonvolatile media on the CMC or stored on the remote share. To initiate a server configuration replication operation, you can apply a stored profile to one or more servers.

NOTE: If a server does not support Dell Lifecycle Controller or the chassis is turned off, you cannot apply a profile to the server.

To apply a profile to one or more server(s):

- 1. Go to the **Server Profiles** page. In the **Save and Apply Profiles** section, select the server or servers for which you want to apply the selected profile.
 - The **Select Profile** drop-down menu is enabled.
 - NOTE: The **Select Profile** drop-down menu displays all available profiles and sorted by type, including those that are on the remote share and SD card.
- 2. From the **Select Profile** drop-down menu, select the profile that you want to apply.
 - The **Apply Profile** option is enabled.
- 3. Click Apply Profile.

A message is displayed that applying a new server profile overwrites the current settings and also restarts the selected servers. You are prompted to confirm if you want to continue the operation.

- NOTE: To perform server cloning operations on servers, the CSIOR option must be enabled for the servers. If CSIOR option is disabled, a warning message is displayed that CSIOR is not enabled for the servers. To complete the blade cloning operation, make sure to enable CSIOR option on the servers.
- **4.** Click **OK** to apply the profile to the selected server.

The selected profile is applied to the server(s) and the server(s) may be restarted immediately, if necessary. For more information, see the CMC Online Help.

Importing Profile

You can import a server profile that is stored on a management station to CMC.

To import a stored profile from CMC:

- In the Server Profiles page, in the Stored Profiles section, click Import Profile.
 The Import Server Profile section is displayed.
- 2. Click Browse to access the profile from the required location and then click Import Profile.

For more information about the fields, see the Online Help.

Exporting Profile

You can export a stored server profile to a specified file folder path on a management station.

To export a stored profile:

- Go to the Server Profiles page. In the Stored Profiles section, select the required profile, and then click Export Copy of Profile.
 - A File Download message is displayed prompting you to open or save the file.
- 2. Click Save or Open to export the profile to the required location.
 - NOTE: If the source profile is on the SD card, a message is displayed indicating that if the profile is exported, then the description is lost. Press **OK** to continue exporting the profile.

A message is displayed prompting you to select the destination of the file:

- Local or Network Share if the source file is on an SD card.
 - NOTE: The Network Share option is enabled and the details are displayed in the Stored Profiles section only if the network share is mounted and is accessible. If the Network Share is not connected, configure the Network Share for the chassis. To configure the Network Share, click Edit in the Stored Profiles section. For more information, see Configuring Network Share Using CMC Web Interface.
- Local or SD Card if the source file is on the Network Share.

For more information about the fields, see the Online Help.

- 3. Select Local, Extended Storage, or Network Share as the destination location based on the options displayed.
 - If you select Local, a dialog box appears allowing you to save the profile to a local directory.
 - If you select Extended Storage or Network Share, a Save Profile dialog box is displayed.
- 4. Click Save Profile to save the profile to the selected location.
- NOTE: The CMC web interface captures the normal server configuration profile (snapshot of the server), which can be used for replication on a target system. However, some configurations such as RAID and identity attributes are not propagated to the new server. For more information on alternate export modes for RAID configurations and identity attributes, see the white paper, Server Cloning with Server Configuration Profiles, at DellTechCenter.com.

Editing Profile

You can edit the name and description of a server profile that is stored on the CMC nonvolatile media (SD Card) or the name of a server profile stored on the remote share.

To edit a stored profile:

- Go to the Server Profiles page. In the Stored Profiles section, select the required profile and then click Edit Profile.
 The Edit Server Profile <Profile Name> section is displayed.
- 2. Edit the profile name and description of the server profile as required and then click Edit Profile.
 - i NOTE: You can edit the profile description only for profiles stored on SD cards.

For more information, see the Online Help.

Deleting Profile

You can delete a server profile that is stored on the CMC nonvolatile media (SD Card) or on the Network Share.

To delete a stored profile:

- In the Server Profiles page, in the Stored Profiles section, select the required profile and then click Delete Profile.
 A warning message is displayed indicating that deleting a profile would permanently delete the selected profile.
- 2. Click **OK** to delete the selected profile.

For more information, see the Online Help.

Viewing Profile Settings

To view Profile settings for a selected server, go to the **Server Profiles** page. In the **Server Profiles** section, click **View** in the **Server Profile** column for the required server. The **View Settings** page is displayed.

For more information about the displayed settings, see the Online Help.

NOTE: The CMC Server Configuration Replication feature retrieves and displays the settings for a specific server, only if the Collect System Inventory on Restart (CSIOR) option is enabled.

To enable CSIOR on:

• 12th generation servers — After restarting the server, when the company logo is displayed, select F2. On the **iDRAC Settings** page, in the left pane, click **Lifecycle Controller**, and then click **CSIOR** to enable the changes.

• 13th generation servers —After restarting the server, when prompted, press F10 to access Dell Lifecycle Controller. Go to the **Hardware Inventory** page by clicking **Hardware Configuration** > **Hardware Inventory**. On the **Hardware Inventory** page, click **Collect System Inventory on Restart**.

Viewing Stored Profile Settings

To view profile settings of the stored server profiles, go to the **Server Profiles** page. In the **Stored Profiles** section, click **View in the View Profile** column for the required server profile. The **View Settings** page is displayed. For more information about the displayed settings, see the *Online Help*.

Viewing Profile Log

To view the profile log, in the **Server Profiles** page, see the **Recent Profile Log** section. This section lists the 10 latest profile log entries directly from server configuration operations. Each log entry displays the severity, the time and date of submission of the server configuration replication operation, and the replication log message description. The log entries are also available in the RAC log. To view the other available entries, click **Go to Profile Log**. The **Profile Log** page is displayed. For more information, see the *Online Help*.

Completion Status And Troubleshooting

To check the completion status of an applied BIOS profile:

- 1. In the left pane, click Chassis Overview > Server Overview > Setup > Profiles.
- 2. On the Server Profiles page, note down the Job ID (JID) of the submitted job from the Recent Profile Log section.
- 3. In the left pane, click Server Overview > Troubleshooting > Lifecycle Controller Jobs. Search for the same JID in the Jobs table. For more information about performing Lifecycle Controller jobs using CMC, see Lifecycle Controller Job Operations.
- 4. Click View Log link to view the results of Lclogview from the iDRAC Lifecycle Controller for the specific server. The results displayed for the completion or failure are similar to the information displayed in the iDRAC Lifecycle Controller log for the specific server.

Quick Deploy of Profiles

The Quick Deploy feature enables you to assign a stored profile to a server slot. Any server supporting server configuration replication that is inserted into a slot is configured using the profile assigned to that slot. You can perform the Quick Deploy action only if the **Action When Server is Inserted** option on the **Deploy iDRAC** page is set to **Server Profile** option or **Quick Deploy and Server Profile** option. Selecting this option allows to apply the server profile assigned when a new server is inserted in the chassis. To go to the **Deploy iDRAC** page, select **Server Overview** > **Setup** > **iDRAC**. Profiles that can be deployed are stored in the SD card or remote share. To set up the profiles for quick deploy, you must have **Chassis Administrator** privileges.



Assigning Server Profiles to Slots

The Server Profiles page enables you to assign server profiles to slots. To assign a profile to the chassis slots:

- 1. On the Server Profiles page, click Profiles for QuickDeploy.
 - The current profile assignments are displayed for the slots in the select boxes contained in the **Assign Profile** column.
 - NOTE: You can perform the Quick Deploy action only if the Action When Server is Inserted option in the Deploy iDRAC page is set to Server Profile or Quick Deploy then Server Profile. Selecting this option allows to apply the server profile assigned when a new server is inserted in the chassis.
- 2. From the drop-down menu, select the profile to assign to the required slot. You can select a profile to apply to multiple slots.
- 3. Click Assign Profile.
 - The profile is assigned to the selected slots.

(i) NOTE:

- A slot that does not have any profile assigned to it is indicated by the term "No Profile Selected" that appears in the select box.
- To remove a profile assignment from one or more slots, select the slots and click **Remove Assignment** A message is displayed warning you that removing a profile from the slot or slots removes the configuration settings in the profile from any server (s) inserted in the slot (s) when **Quick Deploy Profiles** feature is enabled. Click **OK** to remove the profile assignments.
- To remove all profile assignments from a slot, in the drop-down menu, select **No Profile Selected**.
- NOTE: When a profile is deployed to a server using the **Quick Deploy Profiles** feature, the progress and results of the application are retained in the Profile Log.

(i) NOTE:

- If an assigned profile is on the Network Share which is not accessible when a server is inserted in the slot, the LCD displays a message that the assigned profile is not available for Slot <X>.
- The **Network Share** option is enabled and the details are displayed in the **Stored Profiles** section only if the network share is mounted and is accessible. If the Network Share is not connected, configure the Network Share for the chassis. To configure the Network Share, click **Edit** in the **Stored Profiles** section. For more information, see Configuring Network Share Using CMC Web Interface.

Boot Identity Profiles

To access the **Boot Identity Profiles** page in the CMC web interface, in the system tree, go to **Chassis Overview > Server Overview**. Click **Setup > Profiles**. The **Server Profiles** page is displayed. On the **Server Profiles** page, click **Boot Identity Profiles**.

The boot identity profiles contain the NIC or FC settings that are required to boot a server from a SAN target device and unique virtual MAC and WWN. As these are available across multiple chassis through a CIFS or NFS share, you can quickly and remotely move an identity from a non-functional server in a chassis to a spare server located in the same or another chassis and thus enabling it to boot with the operating system and applications of the failed server. The main advantage of this feature is the use of a virtual MAC address pool that is unique and shared across all chassis.

This feature allows you to manage server operations online without physical intervention if the server stops functioning. You can perform the following tasks by using the Boot Identity Profiles feature:

- Initial setup
 - Create a range of virtual MAC addresses. To create a MAC address, you must have Chassis Configuration Administrator and Server Administrator privileges.
 - Save boot identity profile templates and customize the boot identity profiles on the network share by editing and including the SAN boot parameters that are used by each server.
 - Prepare the servers that use initial configuration before applying their Boot Identity profiles.
 - Apply Boot Identity profiles to each server and boot them from SAN.
- Configure one or more spare standby servers for quick recovery.
 - o Prepare the standby servers that use initial configuration before applying their Boot Identity profiles.
- Use the workload of a failed server in a new server by performing the following tasks:
 - Clear the boot identity from the non-functioning server to avoid duplicating the MAC addresses in case the server recovers.
 - o Apply the boot identity of a failed server to a spare standby server.
 - o Boot the server with the new Boot Identity settings to quickly recover the workload.

Saving Boot Identity Profiles

You can save boot identity profiles in the CMC network share. Number of profiles that you can store depends on the availability of MAC addresses. For more information, see *Configuring Network Share Using CMC Web Interface*.

For Emulex Fibre Channel (FC) cards, the **Enable/Disable Boot From SAN** attribute in the Option ROM is disabled by default. Enable the attribute in the Option ROM and apply the boot identify profile to the server for booting from SAN.

To save a profile, perform the following tasks:

- 1. Go to the **Server Profiles** page. In the **Boot Identity Profiles** section, select the server that has the required settings with which you want to generate the profile and select FQDD from the **FQDD** drop-down menu.
- 2. Click Save Identity. The Save Identity section is displayed.
 - NOTE: Boot identity is saved only if the **Network Share** option is enabled and accessible, the details are displayed in the **Stored Profiles** section. If the **Network Share** is not connected, configure the network share for the chassis. To configure the network share, click **Edit** in the **Stored Profiles** section. For more information, see *Configuring Network Share Using CMC Web Interface*.
- 3. In the **Base Profile Name** and **Number of Profiles** fields, enter the profile name and the number of profiles that you want to save.
 - NOTE: While saving a boot identity profile, the standard ASCII extended character set is supported. However, the following special characters are not supported:

4. Select a MAC address for the base profile from the Virtual MAC Address drop-down and click Save Profile.

The number of templates created are based on the number of profiles you specify. CMC communicates with the Lifecycle Controller to get the available server profile settings and store them as a named profile. The format for the name file is — <base profile name> <profile number> <MAC address>. For example: FC630 01 0E00000000000.

A progress indicator indicates that the save operation is in progress. After the action is complete, **Operation Successful** message is displayed.

NOTE: The process to gather the settings occurs in the background. Hence, it may take some time before the new profile is displayed. If the new profile is not displayed, check the profile log for errors.

Applying Boot Identity Profiles

You can apply boot identity profile settings if the boot identity profiles are available as stored profiles on the network share. To initiate a boot identity configuration operation, you can apply a stored profile to a single server.

NOTE: If a server does not support Lifecycle Controller or the chassis is powered off, you cannot apply a profile to the server.

To apply a profile to a server, perform the following tasks:

1. Go to the **Server Profiles** page. In the **Boot Identity profiles** section, select the server on which you want to apply the selected profile.

The **Select Profile** drop-down menu gets enabled.

- NOTE: The Select Profile drop-down menu displays all available profiles that are sorted by type from the network share.
- 2. From the **Select Profile** drop-down menu, select the profile that you want to apply. The **Apply Identity** option is enabled.
- 3. Click Apply Identity.

A warning message is displayed that applying a new identity overwrites the current settings and also reboots the selected server. You are prompted to confirm if you want to continue the operation.

- NOTE: To perform server configuration replication operations on the server, the CSIOR option must be enabled for the servers. If CSIOR option is disabled, a warning message is displayed that CSIOR is not enabled for the server. To complete the server configuration replication operation, enable the CSIOR option on the server.
- $\textbf{4.} \ \ \, \text{Click } \textbf{OK} \text{ to apply the boot identity profile to the selected server}.$

The selected profile is applied to the server and the server is rebooted immediately. For more information, see the CMC Online Help.

NOTE: You can apply a boot identity profile to only one NIC FQDD partition in a server at a time. To apply the same boot identity profile to a NIC FQDD partition in another server, you must clear it from the server where it is first applied.

Clearing Boot Identity Profiles

Before applying a new boot identity profile to a standby server, you can clear the existing boot identity configurations of a selected server by using the **Clear Identity** option available in the CMC web interface.

To clear boot identity profiles:

- Go to the Server Profiles page. In the Boot Identity profiles section, select the server from which you want to clear the boot identity profile.
 - NOTE: This option is enabled only if any of the servers are selected and boot identity profiles are applied to the selected servers.
- 2. Click Clear Identity.
- 3. Click **OK** to clear the boot identity profile from the selected server.

 The clear operation disables the IO Identity and persistence policy of the server. On completion of the clear operation, the server is powered off.

Viewing Stored Boot Identity Profiles

To view the boot identity profiles stored on the network share, go to the **Server Profiles** page. In the **Boot Identity Profiles** > **Stored Profiles** section, select the profile and click **View** in the **View Profile** column. The **View Settings** page is displayed. For more information on the displayed settings, see the *CMC Online Help*.

Importing Boot Identity Profiles

You can import boot identity profiles that are stored on the management station to the network share.

To import a stored profile on to the network share from the management station, perform the following tasks:

- Go to the Server Profiles page. In the Boot Identity Profiles > Stored Profiles section, click Import Profile.
 The Import Profile section is displayed.
- 2. Click **Browse** to access the profile from the required location and then click **Import Profile**. For more information, see the *CMC Online Help*.

Exporting Boot Identity Profiles

You can export a boot identity profiles that are saved on the network share to a specified path on a management station.

To export a stored profile, perform the following tasks:

- Go to the Server Profiles page. In the Boot Identity Profiles > Stored Profiles section, select the required profile and then click Export Profile.
 - A File Download message is displayed prompting you to open or save the file.
- 2. Click Save or Open to export the profile to the required location.

Deleting Boot Identity Profiles

You can delete a boot identity profile that is stored on the network share.

To delete a stored profile, perform the following tasks::

- Go to the Server Profiles page. In the Boot Identity Profiles > Stored Profiles section, select the required profile, and then click Delete Profile.
 - A warning message is displayed indicating that deleting a profile would delete the selected profile permanently.
- 2. Click **OK** to delete the selected profile.
 - For more information, see the CMC Online Help.

Managing Virtual MAC Address Pool

You can create, add, remove, and deactivate MAC addresses by using the **Managing Virtual MAC Address Pool**. You can only use unicast MAC addresses in the Virtual MAC Address Pool. The following MAC address ranges are allowed in CMC.

- 02:00:00:00:00 F2:FF:FF:FF:FF
- 06:00:00:00:00:00 F6:FF:FF:FF:FF
- 0A:00:00:00:00 FA:FF:FF:FF:FF
- 0E:00:00:00:00:00 FE:FF:FF:FF:FF

To view the Manage Virtual MAC Address option by the CMC web interface, in the system tree, go to Chassis Overview > Server Overview. Click Setup > Profiles > Boot Identity Profiles. The Manage Virtual MAC Address Pool section is displayed.

NOTE: The virtual MAC Addresses are managed in the vmacdb.xml file in the network share. A hidden lock file (.vmacdb.lock) is added and removed from the network share to serialize boot identity operations from multiple chassis.

Creating MAC Pool

You can create MAC pool in the network by using the **Manage Virtual MAC Address Pool** option available in the CMC web interface.

NOTE: The Create MAC Pool section is displayed only if the MAC address database (vmacdb.xml) is not available in the network share. In this case, the Add MAC Address and Remove MAC Address options are disabled.

To create a MAC pool:

- 1. Go to the Server Profiles page. In the Boot Identity Profiles > Manage Virtual MAC Address Pool section.
- 2. Enter the starting MAC address of the MAC address pool in the Starting MAC Address field.
- 3. Enter the count of the MAC addresses in the Number of MAC Addresses field.
- 4. Click Create MAC Pool to create the MAC address pool. After the MAC address database is created in the network share, the Manage Virtual MAC Address Pool displays the list and status of the MAC addresses that are stored in the network share. This section now enables you to add or remove MAC addresses from the MAC Address Pool.

Adding MAC Addresses

You can add a range of MAC addresses to the network share by using the **Add MAC Addresses** option available in the CMC web interface.

NOTE: You cannot add a MAC address that exists in the MAC address pool. An error is displayed indicating that the newly added MAC address exists in the pool.

To add MAC addresses to the network share:

- Go to the Server Profiles page. In the Boot Identity Profiles > Manage Virtual MAC Address Pool section, click Add MAC Addresses.
- 2. Enter the starting MAC address of the MAC address pool in the Starting MAC Address field.
- **3.** Enter the count of the MAC addresses that you want to add, in the **Number of MAC Addresses** field. The valid values are from 1 to 3000.
- 4. Click OK to add MAC addresses.

For more information, see the CMC Online Help.

For more information, see the CMC for Dell PowerEdge FX2/FX2s Online Help.

Removing MAC Addresses

You can remove a range of MAC addresses from the network share by using the **Remove MAC Addresses** option available in the CMC web interface.

(i) NOTE: You cannot remove MAC addresses if they are active on the node or are assigned to a profile.

To remove MAC addresses from the network share:

- Go to the Server Profiles page. In the Boot Identity Profiles > Manage Virtual MAC Address Pool section, click Remove MAC Addresses.
- 2. Enter the starting MAC address of the MAC address pool in the Starting MAC Address field.
- 3. Enter the count of the MAC addresses that you want to remove, in the Number of MAC Addresses field.
- 4. Click **OK** to remove MAC addresses.

Deactivating MAC Addresses

You can deactivate MAC addresses that are active by using the **Deactivate MAC Address(es)** option in the CMC web interface.

NOTE: Use the **Deactivate MAC Address(es)** option only if the server is not responding to the **Clear Identity** action or the MAC address is not used in any server.

To remove MAC addresses from the network share:

- Go to the Server Profiles page. In the Boot Identity Profiles > Manage Virtual MAC Address Pool section, select the
 active MAC address(es) that you want to deactivate.
- 2. Click Deactivate MAC Address(es).

Launching iDRAC using Single Sign-On

CMC provides limited management of individual chassis components, such as servers. For complete management of these individual components, CMC provides a launch point for the server's management controller (iDRAC) web-based interface.

A user may be able to launch iDRAC web interface without having to log in a second time, as this feature utilizes single sign-on. Single sign-on policies are:

- A CMC user who has server administrative privilege, is automatically logged in to iDRAC using single sign-on. Once on the
 iDRAC site, this user is automatically granted Administrator privileges. This is true even if the same user does not have an
 account on iDRAC, or if the account does not have the Administrator's privileges.
- A CMC user who does NOT have the server administrative privilege, but has the same account on iDRAC is automatically
 logged in to iDRAC using single sign-on. Once on the iDRAC site, this user is granted the privileges that were created for the
 iDRAC account.
- A CMC user who does not have the server administrative privilege, or the same account on the iDRAC, is not automatically
 logged in to iDRAC using single sign-on. This user is directed to the iDRAC login page when the Launch iDRAC GUI is
 clicked.
- NOTE: The term "the same account" in this context means that the user has the same login name with a matching password for CMC and for iDRAC. The user who has the same login name without a matching password, is considered to have the same account.
- i) NOTE: Users may be prompted to log in to iDRAC (see the third Single Sign-on policy bullet above).
- i NOTE: If the iDRAC network LAN is disabled (LAN Enabled = No), single sign-on is not available.

If you click Launch iDRAC GUI, an error page may be displayed, if:

- the server is removed from the chassis
- the iDRAC IP address is changed
- the iDRAC network connection experiences a problem.

In MCM, while launching the iDRAC web interface from a member chassis, the user credentials of the leader and member chassis must be the same. Else, the current member chassis session is aborted and the member chassis login page is displayed.

Launching iDRAC from Server Status Page

To launch the iDRAC management console for an individual server:

- 1. In the left pane, expand Server Overview. All four servers appear in the expanded Servers Overview list.
- 2. Click the server for which you want to launch the iDRAC Web interface.
- 3. On the Servers Status page, click Launch iDRAC.

The iDRAC Web interface is displayed. For information about the field descriptions, see the Online Help.

Launching iDRAC from Servers Status Page

To launch the iDRAC management console from the Servers Status page:

- 1. In the left pane, click Server Overview.
- 2. On the Servers Status page, click Launch iDRAC for the server you want to launch the iDRAC Web interface.

Launching Remote Console

You can launch a Keyboard-Video-Mouse (KVM) session directly on the server. The remote console feature is supported only when all of the following conditions are met:

- The chassis power is on.
- Servers that support iDRAC7 and iDRAC8.
- The LAN interface on the server is enabled.
- The host system is installed with JRE (Java Runtime Environment) 6 Update 16 or later.
- The browser on host system allows pop-up windows (pop-up blocking is disabled).

Remote Console can also be launched from the iDRAC Web interface. For more details, see the iDRAC User's Guide available at dell.com/support/manuals.

Launching Remote Console from Chassis Health Page

To launch a remote console from the CMC Web interface:

- 1. In the left pane, click Chassis Overview, and then click Properties.
- 2. On the Chassis Health page, click the specified server in the chassis graphic.
- 3. In the Quicklinks section, click the Remote Console link to launch the remote console.

Launching Remote Console from Server Status Page

To launch a remote console for an individual server:

- 1. In the left pane, expand Server Overview. All the four servers appear in the expanded servers' list.
- 2. Click the server for which you want to launch the remote console.
- 3. On the Server Status page, click Launch Remote Console.

Launching Remote Console from Servers Status Page

To launch a remote console from the **Servers Status** page:

- 1. In the left pane, go to Server Overview, and then click Properties > Status. The Servers Status page is displayed.
- 2. Click Launch Remote Console for the required server.

Configuring CMC To Send Alerts

You can set alerts and actions for certain events that occur on the chassis. An event is generated when a device or service's status has changed or an error condition is detected. If an event matches an event filter and you have configured this filter to generate an alert message (email alert or SNMP trap), then an alert is sent to one or more configured destinations such as email address, IP address, or an external server.

To configure CMC to send alerts:

- 1. Enable the Chassis Event Alerts option.
- 2. Optionally, filter the alerts based on category or severity.
- **3.** Configure the email alert or SNMP trap settings.
- 4. Enable chassis event alerts to send an e-mail alert, or SNMP traps to configured destinations.

Topics:

- Enabling Or Disabling Alerts
- Configuring Alert Destinations

Enabling Or Disabling Alerts

To send alerts to configured destinations, you must enable the global alerting option. This property overrides the individual alert setting.

Make sure that the SNMP or email alert destinations are configured to receive the alerts.

Enabling Or Disabling Alerts Using CMC Web Interface

To enable or disable generating alerts:

- 1. In the left pane, click Chassis Overview > Alerts.
- 2. On the Chassis Events page, under the Chassis Alert Enablement section, select the Enable Chassis Event Alerts option to enable, or clear the option to disable the alert.
- 3. To save the settings, click Apply.

Filtering Alerts

You can filter alerts on the basis of category and severity.

Filtering Alerts Using CMC Web Interface

To filter the alerts on the basis of category and severity:

- (i) NOTE: To apply chassis events configuration changes, you must have the Alert Configuration privilege.
- 1. In the left pane, click Chassis Overview > Alerts.
- 2. On the Chassis Events page, under the Alerts Filter section, select one or more of the following categories:
 - System Health
 - Storage
 - Configuration
 - Audit
 - Updates
- **3.** Select one or more of the following severity levels:

- Critical
- Warning
- Informational

The **Monitored Alerts** section displays the results based on the selected category and severity. For information about the field descriptions on the this page, see the *Online Help*.

4. Click Apply.

Setting Event Alerts Using RACADM

To set an event alert, run the eventfilters command. For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support//manuals.

Configuring Alert Destinations

The management station uses Simple Network Management Protocol (SNMP) to receive data from CMC.

You can configure the IPv4 and IPv6 alert destinations, email settings, and SMTP server settings, and test these settings.

Before configuring the email alert or SNMP trap settings, make sure that you have the Chassis Configuration Administrator privilege.

Configuring SNMP Trap Alert Destinations

You can configure the IPv6 or IPv4 addresses to receive the SNMP traps.

Configuring SNMP Trap Alert Destinations Using CMC Web Interface

To configure IPv4 or IPv6 alert destination settings using CMC Web interface:

- 1. In the left pane, click Chassis Overview > Alerts > Trap Settings.
- 2. On the Chassis Event Alert Destinations page, type the following:
 - In the **Destination** field, type a valid IP address. Use the quad-dotted IPv4 format, standard IPv6 address notation, or FQDN. For example: 123.123.123.123 or 2001:db8:85a3::8a2e:370:7334 or dell.com.
 - Choose a format that is consistent with the networking technology or infrastructure. The Test Trap functionality cannot detect incorrect choices based on the current network configuration (example, use of an IPv6 destination in an IPv4-only environment).
 - In the Community String field, enter a valid community name to which the destination management station belongs.
 - This community string differs from the community string on the **Chassis Overview > Network > Services** page. The SNMP traps community string is the community that CMC uses for outbound traps destined to management stations. The community string on the **Chassis Overview > Network > Services** page is the community string that management stations use to query the SNMP daemon on CMC.
 - Under **Enabled**, select the option corresponding to the destination IP to enable the IP address to receive the traps. You can specify up to four IP addresses.
- 3. Click **Apply** to save the settings.
- 4. To test whether the IP address is receiving the SNMP traps, click **Send** in the **Test SNMP Trap** column.

The IP alert destinations are configured.

Configuring SNMP Trap Alert Destinations Using RACADM

To configure IP alert destination using RACADM:

1. Open a serial/Telnet/SSH text console to CMC and log in.

- NOTE: Only one filter mask may be set for both SNMP and email alerting. If you have already selected the filter mask, do not perform task 2, and go to step 3.
- 2. Enable alert generation:

```
racadm config -g cfgAlerting -o cfgAlertingEnable 1
```

- 3. Specify the events filters by running the racadm eventfilters set command.
 - ${f a.}$ To clear all the available alert settings, run the following command: racadm eventfilters set -c cmc.alert.all -n none
 - b. Configure using severity as a parameter. For example, all informational events in storage category are assigned poweroff as action, and email and SNMP as notifications: racadm eventfilters set -c cmc.alert.storage.info -n email, snmp
 - c. Configure using subcategory as a parameter. For example, all configurations under the licensing subcategory in the audit category are assigned poweroff as action and all notifications are enabled: racadm eventfilters set -c cmc.alert.audit.lic -n all
 - d. Configure using subcategory and severity as parameters. For example, all Information events under the licensing subcategory in the audit category are assigned poweroff as action and all notifications are disabled: racadm eventfilters set -c cmc.alert.audit.lic.info -n none
- 4. Enable traps alerts:

```
racadm config -g cfgTraps -o cfgTrapsEnable 1 -i <index>
```

where <index> is a value between 1-4. CMC uses the index number to distinguish up to four configurable destinations for traps alerts. Destinations may be specified as appropriately formatted numeric addresses (IPv6 or IPv4), or Fully-Qualified Domain Names (FQDNs).

5. Specify a destination IP address to receive the traps alert:

```
racadm config -g cfgTraps -o cfgTrapsAlertDestIPAddr <IP address> -i <index>
```

where <IP address> is a valid destination, and <index> is the index value specified in step 4.

6. Specify the community name:

```
racadm config -g cfgTraps -o cfgTrapsCommunityName <community name> -i <index>
```

where <community name> is the SNMP community to which the chassis belongs, and <index> is the index value specified in steps 4 and 5.

You can configure up to four destinations to receive traps alerts. To add more destinations, do the tasks in steps 2–6.

- NOTE: The commands in steps 2-6 overwrites any existing settings configured for the index specified (1-4). To determine whether an index has previously-configured values, type: racadm getconfig -g cfgTraps -i <index>. If the index is configured, values appear for the cfgTrapsAlertDestIPAddr and cfgTrapsCommunityName objects.
- 7. To test an event trap for an alert destination, type:

```
racadm testtrap -i <index>
```

where <index> is a value between 1-4 representing the alert destination you want to test.

If you are not sure about the index number, run the following command:

```
racadm getconfig -g cfgTraps -i <index>
```

Configuring Email Alert Settings

When CMC detects a chassis event, such as an environmental warning or a component failure, it can be configured to send an email alert to one or more email addresses.

Configure the SMTP email server to accept relayed emails from the CMC IP address, a feature which is normally turned off in most mail servers due to security concerns. For instructions to configure in a secure manner, see the documentation that was provided with the SMTP server.

- NOTE: If your mail server is Microsoft Exchange Server 2007, make sure that iDRAC domain name is configured for the mail server to receive the email alerts from iDRAC.
- NOTE: Email alerts support both IPv4 and IPv6 addresses. The DRAC DNS Domain Name must be specified when using IPv6.

If your network has an SMTP server that releases and renews IP address leases periodically, and the addresses are different, then there is a duration when this property setting does not work due to change in the specified SMTP server IP address. In such cases, use the DNS name.

Configuring Email Alert Settings Using CMC Web Interface

To configure the email alert settings using Web interface:

- 1. In the left pane, click Chassis Overview > Alerts > E-mail Alert Settings.
- 2. Specify the SMTP email server settings and the email addresses to receive the alerts. For information about the field descriptions, see the *Online Help*.
- **3.** Click **Apply** to save the settings.
- 4. Click **Send** under **Test E-mail** to send a test email to the specified email alert destination.

Configuring EMail Alert Settings Using RACADM

To send a test email to an email alert destination using RACADM:

- 1. Open a serial/Telnet/SSH text console to CMC and log in.
- 2. Enable alert generation:

```
racadm config -g cfgAlerting -o cfgAlertingEnable 1
```

- NOTE: Only one filter mask may be set by both SNMP and email alerting. If you have already set a filter mask, do not perform the task in step 3.
- **3.** Specify the events for which alerts must be generated:

```
racadm config -g cfgAlerting -o cfgAlertingFilterMask <mask value>
```

where <mask value> is a hexadecimal value between 0x0 and 0xfffffffff and must be expressed with the leading 0x characters. The Event Traps Filter Masks table provides filter masks for each event type. For instructions about calculating the hex value for the filter mask you want to enable, see step 3 in the Configuring SNMP Trap Alert Destinations Using RACADM.

4. Enable email alert generation:

```
racadm config -g cfgEmailAlert -o cfgEmailAlertEnable 1 -i <index>
```

where <index> is a value between 1-4. CMC uses the index number to distinguish up to four destination email addresses that can be configured.

5. Specify a destination email address to receive the email alerts:

```
racadm config -g cfgEmailAlert -o cfgEmailAlertAddress <email address> -i <index>
```

where <email address> is a valid email address, and <index> is the index value you specified in step 4.

6. Specify the name of the person receiving the email alert:

```
racadm config -g cfgEmailAlert -o cfgEmailAlertEmailName <email name> -i <index>
```

where <email name> is the name of the person or group receiving the email alert, and <index> is the index value specified in step 4 and step 5. The email name can contain up to 32 alphanumeric characters, dashes, underscores, and periods. Spaces are not valid.

7. Set up the SMTP host:

 $\verb|racadm| config -g cfgRemoteHosts -o cfgRhostsSmtpServerIpAddr host.domain| \\$

where host.domain is the FQDN.

You can configure up to four destination email addresses to receive email alerts. To add more email addresses, perform tasks in step 2 through 6.

NOTE: The commands in steps 2-6 overwrite any existing settings configured for the index that you specify (1-4). To determine whether an index has previously configured values, type racadm getconfig -g cfgEmailAlert - I <index>. If the index is configured, values appear for the cfgEmailAlertAddress and cfgEmailAlertEmailName objects.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Configuring User Accounts and Privileges

You can set up user accounts with specific privileges (role-based authority) to manage your system with CMC and maintain system security. By default, CMC is configured with a local administrator account. The default user name is rootand the password is calvin. As an administrator, you can set up user accounts to allow other users to access the CMC.

You can set up a maximum of 16 local users, or use directory services such as Microsoft Active Directory or LDAP to set up additional user accounts. Using a directory service provides a central location for managing authorized user accounts.

CMC supports role-based access to users with a set of associated privileges. The roles are administrator, operator, read-only, or none. The role defines the maximum privileges available.

Topics:

- Types of Users
- Modifying Root User Administrator Account Settings
- Configuring Local Users
- Configuring Active Directory Users
- Configuring Generic LDAP Users

Types of Users

There are two types of users:

- CMC users or chassis users
- iDRAC users or server users (since the iDRAC resides on a server)

CMC and iDRAC users can be local or directory service users.

Except where a CMC user has **Server Administrator** privilege, privileges granted to a CMC user are not automatically transferred to the same user on a server, because server users are created independently from CMC users. In other words, CMC Active Directory users and iDRAC Active Directory users reside on two different branches in the Active Directory tree. To create a local server user, the Configure Users must log in to the server directly. The Configure Users cannot create a server user from CMC or vice versa. This rule protects the security and integrity of the servers.

Table 19. User Types

Privilege	Description
CMC Login User	User can log in to CMC and view all the CMC data, but cannot add or modify data or execute commands.
	It is possible for a user to have other privileges without the CMC Login User privilege. This feature is useful when a user is temporarily not allowed to login. When that user's CMC Login User privilege is restored, the user retains all the other privileges previously granted.
Chassis Configuration Administrator	 User can add or change data that: Identifies the chassis, such as chassis name and chassis location. Is assigned specifically to the chassis, such as IP mode (static or DHCP), static IP address, static gateway, and static subnet mask. Provides services to the chassis, such as date and time, firmware update, and CMC reset. Is associated with the chassis, such as slot name and slot priority. Although these properties apply to the servers, they are strictly chassis properties relating to the slots rather than the servers themselves. For this reason, slot names and slot priorities can be added or changed whether or not servers are present in the slots. When a server is moved to a different chassis, it inherits the slot name and priority assigned to the slot it occupies in the new chassis. The previous slot name and priority remain with the previous chassis.

Table 19. User Types (continued)

Privilege	Description		
	NOTE: CMC users with the Chassis Configuration Administrator privilege can configure power settings. However, the Chassis Control Administrator privilege is required to perform chassis power operations, including power on, power off, and power cycle.		
User Configuration Administrator	 User can: Add a new user. Change the password of a user. Change the privileges of a user. Enable or disable the login privilege of a user but retain the name and other privileges of the rin the database. 		
Clear Logs Administrator	User can clear the hardware log and CMC log.		
Chassis Control Administrator (Power Commands)	CMC users with the Chassis Power Administrator privilege can perform all power-related operations. They can control chassis power operations, including power on, power off, and power cycle. (i) NOTE: To configure power settings, the Chassis Configuration Administrator privilege is needed.		
Server Administrator	This is a blanket privilege, granting a CMC user all rights to perform any operation on any servers present in the chassis.		
	When a user with Server Administrator privilege issues an action to be performed on a server, the CMC firmware sends the command to the targeted server without checking the privileges of a user on the server. In other words, the Server Administrator privilege overrides any lack of administrator privileges on the server.		
	Without the Server Administrator privilege, a user created on the chassis can only execute a command on a server when all of the following conditions are true: The same user name exists on the server. The same user name must have the same password on the server. The user must have the privilege to execute the command.		
	When a CMC user who does not have Server Administrator privilege issues an action to be performed on a server, CMC sends a command to the targeted server with the user's login name and password. If the user does not exist on the server, or if the password does not match, the user is denied the ability to perform the action.		
	If the user exists on the target server and the password matches, the server responds with the privileges of which the user was granted on the server. Based on the privileges responding from the server, CMC firmware decides if the user has the right to perform the action.		
	Listed below are the privileges and the actions on the server to which the Server Administrator is entitled. These rights are applied only when the chassis user does not have the Server Administrative privilege on the chassis.		
	Server Configuration Administrator: Set IP address Set gateway Set subnet mask Set first boot device		
	Configure Users: • Set iDRAC root password • iDRAC reset		
	Server Control Administrator: Power on Power off Power cycle Graceful shutdown		

Table 19. User Types (continued)

Privilege	Description	
	Server Reboot	
Test Alert User	User can send test alert messages.	
Debug Command Administrator	User can execute system diagnostic commands.	
Fabric A Administrator	User can set and configure the Fabric A IOM.	
Fabric B Administrator	User can set and configure the Fabric B, which corresponds to the first mezzanine card in the servers and is connected to the fabric B circuitry in the shared PCle subsystem in the main board.	
Fabric C Administrator	User can set and configure the Fabric C, which corresponds to the second mezzanine card in the servers and is connected to the fabric C circuitry in the shared PCle subsystem in the main board.	

The CMC user groups provide a series of user groups that have preassigned user privileges.

NOTE: If you select Administrator, Power User, or Guest User, and then add or remove a privilege from the pre-defined set, the CMC Group automatically changes to Custom.

Table 20. CMC Group Privileges

User Group	Privileges Granted	
Administrator	 CMC Login User Chassis Configuration Administrator User Configuration Administrator Clear Logs Administrator Server Administrator Test Alert User Debug Command Administrator Fabric A Administrator 	
Power User	 Login Clear Logs Administrator Chassis Control Administrator (Power commands) Server Administrator Test Alert User Fabric A Administrator 	
Guest User	Login	
Custom	Select any combination of the following permissions: CMC Login User Chassis Configuration Administrator User Configuration Administrator Clear Logs Administrator Chassis Control Administrator (Power commands) Server Administrator Test Alert User Debug Command Administrator Fabric A Administrator	
None	No assigned permissions	

Table 21. Comparison of Privileges Between CMC Administrators, Power Users, and Guest Users

Privilege Set	Administrator	Power User	Guest User
	Permissions	Permissions	Permissions
CMC Login User	Yes	Yes	Yes

Table 21. Comparison of Privileges Between CMC Administrators, Power Users, and Guest Users (continued)

Privilege Set	Administrator Permissions	Power User Permissions	Guest User Permissions
Chassis Configuration Administrator	Yes	No	No
User Configuration Administrator	Yes	No	No
Clear Logs Administrator	Yes	Yes	No
Chassis Control Administrator (Power commands)	Yes	Yes	No
Server Administrator	Yes	Yes	No
Test Alert User	Yes	Yes	No
Debug Command Administrator	Yes	No	No
Fabric A Administrator	Yes	Yes	No

Modifying Root User Administrator Account Settings

For added security, it is strongly recommended that you change the default password of the root (User 1) account. The root account is the default administrative account that is shipped with CMC.

To change the default password for the root account:

- 1. In the left pane, click Chassis Overview, and then click User Authentication.
- 2. On the Users page, in the User ID column, click 1.
 - NOTE: The user ID 1 is the root user account that is shipped by default with CMC. This cannot be changed.
- 3. On the User Configuration page, select the Change Password option.
- 4. Type the new password in the Password field, and then type the same password in Confirm Password.
- 5. Click Apply. The password is changed for the 1 user ID.

Configuring Local Users

You can configure up to 16 local users in CMC with specific access privileges. Before you create a CMC local user, verify if any current users exist. You can set user names, passwords, and roles with the privileges for these users. The user names and passwords can be changed using any of the CMC-secured interfaces such as, web interface, RACADM, and WS-MAN.

Configuring Local Users Using CMC Web Interface

(i) NOTE: You must have Configure Users permission to create a CMC user.

To add and configure local CMC users:

- 1. In the left pane, click Chassis Overview, and then click User Authentication.
- 2. On the Local Users page, in the User ID column, click a user ID number. The User Configuration page is displayed.
 - (i) NOTE: User ID 1 is the root user account that is shipped by default with a CMC. This cannot be changed.
- 3. Enable the user ID and specify the user name, password, and access privileges for the user. For more information about the options, see the *Online Help*.
- 4. Click Apply. The user is created with appropriate privileges.

Configure Local Users Using RACADM

(i) NOTE: You must be logged in as a root user to execute RACADM commands on a remote Linux system.

You can configure up to 16 users in the CMC property database. Before you manually enable a CMC user, verify if any current users exist.

If you are configuring a new CMC or if you have used the racadm racresetcfg command, the only current user is root with the password calvin. The racresetcfg subcommand resets all configuration parameters to the default values. Any earlier changes are lost.

i NOTE: Users can be enabled and disabled over time, and disabling a user does not delete the user from the database.

To verify if a user exists, open a Telnet/SSH text console to the CMC, log in, and then type the following command once for each index of 1–16:

```
racadm getconfig -g cfgUserAdmin -i <index>
```

NOTE: You can also type racadm getconfig -f <myfile.cfg> and view or edit the myfile.cfg file, which includes all the CMC configuration parameters.

Several parameters and object IDs are displayed with their current values. Two objects of importance are:

```
# cfgUserAdminIndex=XX
cfgUserAdminUserName=
```

If the cfgUserAdminUserName object has no value, that index number, which is indicated by the cfgUserAdminIndex object, is available for use. If a name is displayed after the "=", that index is taken by that user name.

When you manually enable or disable a user with the racadm config subcommand, you must specify the index with the -i option.

The "#" character in the command objects indicates that it is a read-only object. Also, if you use the racadm config -f racadm.cfg command to specify any number of groups/objects to write, the index cannot be specified. A new user is added to the first available index. This behavior allows more flexibility in configuring a second CMC with the same settings as the main CMC.

Adding CMC User Using RACADM

To add a new user to the CMC configuration:

- 1. Set the user name.
- 2. Set the password.
- 3. Set the user privileges. For information about user privileges, see Types of Users.
- 4. Enable the user.

Example:

The following example describes how to add a new user named "John" with a "123456" password and login privileges to the CMC.

```
racadm config -g cfgUserAdmin -o cfgUserAdminUserName
-i 2
john
racadm config -g cfgUserAdmin -o cfgUserAdminPassword
-i 2
123456
racadm config -g cfgUserAdmin -i 2 -o
cfgUserAdminPrivilege
0x00000001
racadm config -g cfgUserAdmin -i 2 -o
cfgUserAdminEnable 1
```

NOTE: For a list of valid bit mask values for specific user privileges, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide. The default privilege value is 0, which indicates that the privileges of a user are not enabled.

To verify that the user was successfully added with the correct privileges, run the following command:

```
racadm getconfig -g cfgUserAdmin -i 2
```

For more information about the RACADM commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/support/manuals**.

Disabling CMC User

When using RACADM, users must be disabled manually on an individual-basis. Users cannot be deleted using a configuration file.

To delete a CMC user, the command syntax is:

```
racadm config -g cfgUserAdmin -o cfgUserAdminUserName
-i <index>""
racadm config -g cfgUserAdmin -i 2 -o
cfgUserAdminPrivilege 0x0
```

A null string of double quotation marks ("") instructs CMC to remove the user configuration at the specified index, and then reset the user configuration to the factory default values.

Enabling CMC User With Permissions

To enable a user with specific administrative permissions (role-based authority):

1. Locate an available user index using the command syntax:

```
racadm getconfig -g cfgUserAdmin -i <index>
```

2. Type the following commands with the new user name and password.

```
racadm config -g cfgUserAdmin -o
cfgUserAdminPrivilege -i <index> <user privilege bitmask value>
```

NOTE: For a list of valid bit mask values for specific user privileges, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals. The default privilege value is 0, which indicates the user does not has any privileges enabled.

Configuring Active Directory Users

If your company uses the Microsoft Active Directory software, you can configure the software to provide access to CMC, allowing you to add and control CMC user privileges to your existing users in your directory service. This is a licensed feature.

- NOTE: On the following Operating Systems, you can recognize the users of CMC users by using Active Directory.
 - Microsoft Windows 2000
 - Microsoft Windows Server 2003
 - Microsoft Windows Server 2008

You can configure user authentication through Active Directory to log in to the CMC. You can also provide role-based authority, which enables an administrator to configure specific privileges for each user.

Supported Active Directory Authentication Mechanisms

You can use Active Directory to define CMC user access using two methods:

- Standard schema solution that uses Microsoft's default Active Directory group objects only.
- Extended schema solution that has customized Active Directory objects provided by Dell. All the access control objects are maintained in Active Directory. It provides maximum flexibility to configure user access on different CMCs with varying privilege levels.

Standard Schema Active Directory Overview

As shown in the following figure, using standard schema for Active Directory integration requires configuration on both Active Directory and CMC.

In Active Directory, a standard group object is used as a role group. A user who has CMC access is a member of the role group. To give this user access to a specific CMC card, the role group name and its domain name need to be configured on the specific CMC card. The role and the privilege level is defined on each CMC card and not in the Active Directory. You can configure up to five role groups in each CMC. The following table shows the default role group privileges.

Table 22. : Default Role Group Privileges

Role Group	Default Privilege Level	Permissions Granted	Bit Mask
1	None	 CMC Login User Chassis Configuration Administrator User Configuration Administrator Clear Logs Administrator Chassis Control Administrator (Power Commands) Server Administrator Test Alert User Debug Command Administrator Fabric A Administrator 	0x00000fff
2	None	 CMC Login User Clear Logs Administrator Chassis Control Administrator (Power Commands) Server Administrator Test Alert User Fabric A Administrator 	0x00000ed9
3	None	CMC Login User	0x00000001
4	None	No assigned permissions	0x0000000
5	None	No assigned permissions	0x0000000

(i) NOTE: The Bit Mask values are used only when setting Standard Schema with the RACADM.

i NOTE: For more information about user privileges, see Types of Users.

Configuring Standard Schema Active Directory

To configure CMC for an Active Directory login access:

- 1. On an Active Directory server (domain controller), open Active Directory Users and Computers Snap-in.
- 2. Using the CMC Web interface or RACADM:
 - a. Create a group or select an existing group.
 - **b.** Configure the role privileges.
- 3. Add the Active Directory user as a member of the Active Directory group to access CMC.

Configuring Active Directory With Standard Schema Using CMC Web Interface

- NOTE: For information about the various fields, see the CMC Online Help.
- In the left pane, go to Chassis Overview, and then click User Authentication > Directory Services. The Directory Services page is displayed.
- 2. Select Microsoft Active Directory (Standard Schema). The settings to be configured for standard schema is displayed on the same page.
- 3. In the **Common Settings** section, specify the following:
 - Select Enable Active Directory and enter the timeout value for Active Directory in the AD Timeout field.
 - To obtain the Active Directory Domain Controllers from a DNS lookup select **Look Up Domain Controllers with DNS**, and then select one of the following:
 - User Domain from Login to perform the DNS lookup with the domain name of the login user.
 - Specify a Domain enter the domain name to use for the DNS lookup
 - To enable CMC to use the specified Active Directory Domain Controller server addresses, select Specify Domain
 Controller Addresses. These server addresses are the addresses of the domain controllers where the user accounts
 and the role groups are located.
- 4. Click Apply to save the settings.
 - NOTE: You must apply the settings before continuing. If you do not apply the settings, the settings are lost when you navigate to the next page.
- 5. In the Standard Schema Role Groups section, click a Role Group. The Configure Role Group page is displayed.
- 6. Specify the group name, domain, and privileges for a role group.
- 7. Click Apply to save the role group settings and then click Go Back To Configuration page.
- 8. If you have enabled certificate validation, you must upload the domain forest root certificate authority-signed certificate to CMC. In the **Manage Certificates** section, type the file path of the certificate or browse to the certificate file. Click **Upload** to upload the file to CMC.
 - NOTE: The File Path value displays the relative file path of the certificate you are uploading. You must type the absolute file path, which includes the full path and the complete file name and file extension.

The SSL certificates for the domain controllers must be signed by the root certificate authority-signed certificate. The root certificate authority-signed certificate must be available on the management station accessing CMC.

- 9. If you have enabled Single Sign-On (SSO), in the **Kerberos Keytab** section, click **Browse**, specify the keytab file and click **Upload**. When the upload is complete, a message is displayed indicating a successful or failed upload.
- 10. Click Apply. The CMC Web server automatically restarts after you click Apply.
- 11. Log out and then log in to CMC to complete the CMC Active Directory configuration.
- 12. Select Chassis in the system tree, and navigate to the Network tab. The Network Configuration page is displayed.
- 13. Under Network Settings, if Use DHCP (for CMC Network Interface IP Address) is selected, select Use DHCP to obtain DNS server address.

To manually enter a DNS server IP address, clear **Use DHCP to obtain DNS server addresses** and type the primary and alternate DNS server IP addresses.

14. Click Apply Changes.

The CMC Standard Schema Active Directory feature configuration is complete.

Configuring Active Directory With Standard Schema Using RACADM

At the RACADM command prompt, run the following commands:

Using config command:

```
racadm config -g cfgActiveDirectory -o cfgADEnable 1
racadm config -g cfgActiveDirectory -o cfgADType 2
racadm config -g cfgStandardSchema -i <index> -o cfgSSADRoleGroupName <common name of
the role group>
```

racadm config -g cfgStandardSchema -i <index> -o cfgSSADRoleGroupDomain <fully qualified domain name> racadm config -g cfgStandardSchema -i <index> -o cfgSSADRoleGroupPrivilege <Bit Mask Value for specific RoleGroup permissions>

racadm config -g cfgActiveDirectory -o cfgADDomainController1 <fully qualified domain name or IP address of the domain controller>
racadm config -g cfgActiveDirectory -o cfgADDomainController2 <fully qualified domain name or IP address of the domain controller>
racadm config -g cfgActiveDirectory -o cfgADDomainController3 <fully qualified domain name or IP address of the domain controller>

- NOTE: Enter the FQDN of the domain controller, not the FQDN of the domain. For example, enter servername.dell.com instead of dell.com.
- (i) NOTE:

At least one of the three addresses is required to be configured. CMC attempts to connect to each of the configured addresses one-by-one until it makes a successful connection. With Standard Schema, these are the addresses of the domain controllers where the user accounts and the role groups are located.

racadm config -g cfgActiveDirectory -o cfgADGlobalCatalog1 <fully qualified domain name or IP address of the domain controller> racadm config -g cfgActiveDirectory -o cfgADGlobalCatalog2 <fully qualified domain name or IP address of the domain controller> racadm config -g cfgActiveDirectory -o cfgADGlobalCatalog3 <fully qualified domain name or IP address of the domain controller>

- NOTE: The Global Catalog server is only required for standard schema when the user accounts and role groups are in different domains. In multiple domain case, only the Universal Group can be used.
- NOTE: The FQDN or IP address that you specify in this field should match the Subject or Subject Alternative Name field of your domain controller certificate if you have certificate validation enabled.

If you want to disable the certificate validation during the SSL handshake, run the following RACADM command:

• Using the config command: racadm config -g cfgActiveDirectory -o cfgADCertValidationEnable 0 In this case, you do not have to upload the Certificate Authority (CA) certificate.

To enforce the certificate validation during SSL handshake (optional):

• Using the config command: racadm config -g cfgActiveDirectory -o cfgADCertValidationEnable 1 In this case, you must upload the CA certificate using the following RACADM command:

racadm sslcertupload -t 0x2 -f <ADS root CA certificate>

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the Global Catalog FQDN. Make sure that DNS is correctly configured.

Extended Schema Active Directory Overview

Using the extended schema solution requires the Active Directory schema extension.

Active Directory Schema Extensions

The Active Directory data is a distributed database of *attributes* and *classes*. The Active Directory schema includes the rules that determine the type of data that can be added or included in the database. One example of a class that is stored in the database is the user class. Some example user class attributes are user's first name, last name, phone number, and so on.

You can extend the Active Directory database by adding your own unique *attributes* and *classes* for specific requirements. Dell has extended the schema to include the necessary changes to support remote management authentication and authorization using Active Directory.

Each attribute or class that is added to an existing Active Directory Schema must be defined with a unique ID. To maintain unique IDs across the industry, Microsoft maintains a database of Active Directory Object Identifiers (OIDs), so that when

companies add extensions to the schema, they can be guaranteed to be unique and not to conflict with each other. To extend the schema in Microsoft's Active Directory, Dell received unique OIDs, unique name extensions, and uniquely linked attribute IDs for the attributes and classes that are added into the directory service.

• Dell extension: dell

• Dell base OID: 1.2.840.113556.1.8000.1280

• RAC LinkID range: 12070 to 12079

Overview of Schema Extensions

Dell has extended the schema to include an Association, Device, and Privilege property. The Association property is used to link together the users or groups with a specific set of privileges to one or more RAC devices. This model provides an administrator maximum flexibility over the different combinations of users, RAC privileges, and RAC devices on the network without much complexity.

When there are two CMCs on the network that you want to integrate with Active Directory for authentication and authorization, create at least one association object and one RAC device object for each CMC. You can create multiple association objects, and each association object can be linked to as many users, groups of users, or RAC device objects as required. The users and RAC device objects can be members of any domain in the enterprise.

However, each association object can be linked (or, may link users, groups of users, or RAC device objects) to only one privilege object. This example allows an administrator to control each user's privileges on specific CMCs.

The RAC device object is the link to RAC firmware for querying Active Directory for authentication and authorization. When a RAC is added to the network, the administrator must configure the RAC and its device object with its Active Directory name so that users can perform authentication and authorization with Active Directory. Additionally, the administrator must add the RAC to at least one association object for users to authenticate.

i NOTE: The RAC privilege object applies to CMC.

You can create as many or as few association objects as required. However, you must create at least one Association Object, and you must have one RAC device object for each RAC (CMC) on the network that you want to integrate with Active Directory.

The Association Object allows as many or as few users and/or groups as well as RAC Device Objects. However, the Association Object only includes one Privilege Object per Association Object. The Association Object connects the *Users* who have *Privileges* on RACs (CMCs).

Additionally, you can configure Active Directory objects in a single domain or in multiple domains. For example, you have two CMCs (RAC1 and RAC2) and three existing Active Directory users (user1, user2, and user3). You want to give user1 and user2 an administrator privilege to both CMCs and give user3 a login privilege to the RAC2 card.

When adding Universal Groups from separate domains, create an Association Object with Universal Scope. The Default Association objects created by the Dell Schema Extender Utility are Domain Local Groups and does not work with Universal Groups from other domains.

To configure the objects for the single domain scenario:

- 1. Create two Association Objects.
- 2. Create two RAC Device Objects, RAC1 and RAC2, to represent the two CMCs.
- 3. Create two Privilege Objects, Priv1 and Priv2, in which Priv1 has all privileges (administrator) and Priv2 has login privilege.
- 4. Group user1 and user2 into Group1.
- 5. Add Group1 as Members in Association Object 1 (A01), Priv1 as Privilege Objects in A01, and RAC1, RAC2 as RAC Devices in A01.
- 6. Add User3 as Members in Association Object 2 (A02), Priv2 as Privilege Objects in A02, and RAC2 as RAC Devices in A02.

To configure the objects for the multiple domain scenario:

- 1. Make that the domain forest function is in Native or Windows 2003 mode.
- 2. Create two Association Objects, A01 (of Universal scope) and A02, in any domain. The figure Setting Up Active Directory Objects in Multiple Domains shows the objects in Domain2.
- 3. Create two RAC Device Objects, RAC1 and RAC2, to represent the two CMCs.
- 4. Create two Privilege Objects, Priv1 and Priv2, in which Priv1 has all privileges (administrator) and Priv2 has login privilege.
- **5.** Group user1 and user2 into Group1. The group scope of Group1 must be Universal.
- 6. Add Group1 as Members in Association Object 1 (A01), Priv1 as Privilege Objects in A01, and RAC1, RAC2 as RAC Devices in A01
- 7. Add User3 as Members in Association Object 2 (A02), Priv2 as Privilege Objects in A02, and RAC2 as RAC Devices in A02.

Configuring Extended Schema Active Directory

To configure Active Directory to access CMC:

- 1. Extend the Active Directory schema.
- 2. Extend the Active Directory Users and Computers Snap-in.
- **3.** Add CMC users and their privileges to Active Directory.
- 4. Enable SSL on each of your domain controllers.
- 5. Configure CMC Active Directory properties using CMC web interface or RACADM.

Extending Active Directory Schema

Extending your Active Directory schema adds a Dell organizational unit, schema classes and attributes, and example privileges and association objects to the Active Directory schema. Before you extend the schema, make sure that you have Schema Admin privileges on the Schema Master Flexible Single Master Operation (FSMO) Role Owner of the domain forest.

You can extend your schema using one of the following methods:

- Dell Schema Extender utility
- LDIF script file

If you use the LDIF script file, the Dell organizational unit is not added to the schema.

The LDIF files and Dell Schema Extender are located on your *Dell Systems Management Tools and Documentation* DVD in the following respective directories:

- DVDdrive:\SYSMGMT\ManagementStation\support\OMActiveDirectory_Tools\Remote_Management_Adv anced\LDIF_Files
- <DVDdrive>:\SYSMGMT\ManagementStation\support\OMActiveDirector y_Tools\Remote_Management_Advanced\Schema Extender

To use the LDIF files, see the instructions in the Release Notes included in the LDIF Files directory.

You can copy and run the Schema Extender or LDIF files from any location.

Using Dell Schema Extender

CAUTION: The Dell Schema Extender uses the SchemaExtenderOem.ini file. To make sure that the Dell Schema Extender utility functions properly, do not modify the name of this file.

- 1. In the Welcome screen, click Next.
- 2. Read and understand the warning and click Next.
- 3. Select Use Current Log In Credentials or enter a user name and password with schema administrator rights.
- 4. Click Next to run the Dell Schema Extender.
- 5. Click Finish.

The schema is extended. To verify the schema extension, use the MMC and the Active Directory Schema Snap-in to verify that the classes and attributes exist. For more information on classes and attributes, see Classes and Attributes. See the Microsoft documentation for details about using the MMC and the Active Directory Schema Snap-in.

Classes And Attributes

Table 23. Class Definitions for Classes Added to the Active Directory Schema

Class Name	Assigned Object Identification Number (OID)
delliDRACDevice	1.2.840.113556.1.8000.1280.1.7.1.1
delliDRACAssociation	1.2.840.113556.1.8000.1280.1.7.1.2
dellRAC4Privileges	1.2.840.113556.1.8000.1280.1.1.1.3
dellPrivileges	1.2.840.113556.1.8000.1280.1.1.1.4
dellProduct	1.2.840.113556.1.8000.1280.1.1.1.5

Table 24. dellRacDevice Class

OID	1.2.840.113556.1.8000.1280.1.7.1.1
Description	Represents the Dell RAC device. The RAC must be configured as delliDRACDevice in Active Directory. This configuration enables CMC to send Lightweight Directory Access Protocol (LDAP) queries to Active Directory.
Class Type	Structural Class
SuperClasses	dellProduct
Attributes	dellSchemaVersion
	dellRacType

Table 25. delliDRACAssociationObject Class

OID	1.2.840.113556.1.8000.1280.1.7.1.2	
Description	Represents the Dell Association Object. The Association Object provides the connection between the users and the devices.	
Class Type	Structural Class	
SuperClasses	Group	
Attributes	dellProductMembers dellPrivilegeMember	

Table 26. dellRAC4Privileges Class

OID	1.2.840.113556.1.8000.1280.1.1.1.3
Description	Defines the privileges (Authorization Rights) for CMC device.
Class Type	Auxiliary Class
SuperClasses	None
Attributes	dellIsLoginUser
	dellIsCardConfigAdmin
	dellIsUserConfigAdmin
	dellIsLogClearAdmin
	dellIsServerResetUser
	dellIsTestAlertUser
	dellIsDebugCommandAdmin
	dellPermissionMask1
	dellPermissionMask2

Table 27. dellPrivileges Class

OID	1.2.840.113556.1.8000.1280.1.1.1.4	
Description	ed as a container Class for the Dell Privileges (Authorization Rights).	
Class Type	Structural Class	
SuperClasses	User	
Attributes	dellRAC4Privileges	

Table 28. dellProduct Class

OID	1.2.840.113556.1.8000.1280.1.1.1.5
Description	The main class from which all Dell products are derived.

Table 28. dellProduct Class (continued)

OID	1.2.840.113556.1.8000.1280.1.1.1.5
Class Type	Structural Class
SuperClasses	Computer
Attributes	dellAssociationMembers

Table 29. List of Attributes Added to the Active Directory Schema

Assigned OID/Syntax Object Identifier	Single Valued
Attribute: dellPrivilegeMember	FALSE
Description : List of dellPrivilege objects that belong to this attribute.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.1	
Distinguished Name : (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
Attribute: dellProductMembers	FALSE
Description : List of dellRacDevices objects that belong to this role. This attribute is the forward link to the dellAssociationMembers backward link.	
Link ID : 12070	
OID : 1.2.840.113556.1.8000.1280.1.1.2.2	
Distinguished Name : (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
Attribute: dellIsCardConfigAdmin	TRUE
Description : TRUE if the user has Card Configuration rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.4	
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: dellIsLoginUser	TRUE
Description : TRUE if the user has Login rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.3	
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: dellIsUserConfigAdmin	TRUE
Description : TRUE if the user has User Configuration Administrator rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.5	
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: delIsLogClearAdmin	TRUE
Description : TRUE if the user has Clear Logs Administrator rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.6	
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: dellIsServerResetUser	TRUE
Description : TRUE if the user has Server Reset rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.7	
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: dellIsTestAlertUser	TRUE
Description : TRUE if the user has Test Alert User rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.10	

Table 29. List of Attributes Added to the Active Directory Schema (continued)

Assigned OID/Syntax Object Identifier	Single Valued
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: dellIsDebugCommandAdmin	TRUE
Description : TRUE if the user has Debug Command Admin rights on the device.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.11	
Boolean (LDAPTYPE_BOOLEAN 1.3.6.1.4.1.1466.115.121.1.7)	
Attribute: dellSchemaVersion	TRUE
Description : The Current Schema Version is used to update the schema.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.12	
Case Ignore String(LDAPTYPE_CASEIGNORESTRING 1.2.840.113556.1.4.905)	
Attribute: dellRacType	TRUE
Description : This attribute is the Current Rac Type for the dellRacDevice object and the backward link to the dellAssociationObjectMembers forward link.	
OID : 1.2.840.113556.1.8000.1280.1.1.2.13	
Case Ignore String(LDAPTYPE_CASEIGNORESTRING 1.2.840.113556.1.4.905)	
Attribute: dellAssociationMembers	FALSE
Description : List of dellAssociationObjectMembers that belong to this Product. This attribute is the backward link to the dellProductMembers Linked attribute.	
Link ID : 12071	
OID : 1.2.840.113556.1.8000.1280.1.1.2.14	
Distinguished Name (LDAPTYPE_DN 1.3.6.1.4.1.1466.115.121.1.12)	
Attribute: dellPermissionsMask1	•
OID : 1.2.840.113556.1.8000.1280.1.6.2.1 Integer (LDAPTYPE_INTEGER)	
Attribute: dellPermissionsMask2	
OID : 1.2.840.113556.1.8000.1280.1.6.2.2 Integer (LDAPTYPE_INTEGER)	

Installing Dell Extension to the Active Directory Users and Computers Snap-In

When you extend the schema in Active Directory, you must also extend the Active Directory Users and Computers Snap-in so the administrator can manage RAC (CMC) devices, users and user groups, RAC associations, and RAC privileges.

When you install your systems management software using the *Dell Systems Management Tools and Documentation* DVD, you can extend the Snap-in by selecting the **Active Directory Users and Computers Snap-in** option during the installation procedure. See the *Dell OpenManage Software Quick Installation Guide* for additional instructions about installing systems management software. For 64-bit Windows Operating Systems, the Snap-in installer is located at <DVDdrive>: \SYSMGMT\ManagementStation\support\OMActiveDirect ory SnapIn64

For more information about the Active Directory Users and Computers Snap-in, see Microsoft documentation.

Adding CMC Users And Privileges To Active Directory

Using the Dell-extended Active Directory Users and Computers Snap-in, you can add CMC users and privileges by creating RAC device, association, and privilege objects. To add each object, perform the following:

- Create a RAC device Object
- Create a Privilege Object

- Create an Association Object
- Add objects to an Association Object

Creating RAC Device Object

To create RAC device object:

- 1. In the MMC Console Root window, right-click a container.
- 2. Select New > Dell Remote Management Object Advanced.
- **3.** On the **New Object** page, type a name for the new object. The name must be identical to the CMC name that you type in the Configuring Active Directory With Standard Schema Using Web Interface.
- 4. Select RAC Device Object and click OK.

Creating Privilege Object

To create a privilege object:

- i NOTE: You must create a privilege object in the same domain as the related association object.
- 1. In the MMC Console Root window, right-click a container.
- 2. Select New > Dell Remote Management Object Advanced.
- 3. On the **New Object** page, type a name for the new object.
- 4. Select Privilege Object and click OK.
- 5. Right-click the privilege object that you created, and then select **Properties**.
- Click the RAC Privileges tab and assign the privileges for the user or group. For more information about CMC user privileges, see Types of Users.

Creating Association Object

The Association Object is derived from a Group and must contain a Group Type. The Association Scope specifies the Security Group Type for the Association Object. When you create an Association Object, choose the Association Scope that applies to the type of objects you intend to add. For example, if you select Universal, the association objects are only available when the Active Directory Domain is functioning in Native Mode or above.

To create association object:

- 1. In the Console Root (MMC) window, right-click a container.
- 2. Select New > Dell Remote Management Object Advanced.
- 3. On the New Object page, type a name for the new object and select Association Object.
- 4. Select the scope for the Association Object and click OK.

Adding Objects To Association Object

Using the **Association Object Properties** window, you can associate users or user groups, privilege objects, and RAC devices or RAC device groups. If your system is running on Microsoft Windows 2000 operating system or later version, use Universal Groups to span domains with your user or RAC objects.

You can add groups of Users and RAC devices.

Adding Users Or User Groups

To add users or user groups:

- 1. Right-click the Association Object and select Properties.
- 2. Select the Users tab and click Add.
- 3. Enter the user or user group name and click **OK**.

Adding Privileges

To add privileges:

- 1. Select the Privileges Object tab and click Add.
- 2. Enter the privilege object name and click OK.

Click the **Privilege Object** tab to add the privilege object to the association that defines the user's or user group's privileges when authenticating to an RAC device. Only one privilege object can be added to an Association Object.

Adding RAC Devices Or RAC Device Groups

To add RAC devices or RAC device groups:

- 1. Select the Products tab and click Add.
- 2. Enter RAC devices or RAC device group name and click OK.
- 3. In the Properties window, click Apply and click OK.

Click the **Products** tab to add one or more RAC devices to the association. The associated devices specify the RAC devices connected to the network that are available for the defined users or user groups. Multiple RAC devices can be added to an Association Object.

Configuring Active Directory With Extended Schema Using CMC Web Interface

To configure Active Directory with extended schema using CMC web interface:

- i NOTE: For information about the various fields, see the Online Help.
- 1. In the left pane, click, Chassis Overview > User Authentication > Chassis Overview > Directory Services.
- Select Microsoft Active Directory (Extended Schema).
 The settings to be configured for extended schema is displayed on the same page.
- 3. In the Common Settings section, specify the following:
 - Select Enable Active Directory and enter the timeout value for Active Directory in the AD Timeout field.
 - To obtain the Active Directory Domain Controllers from a DNS lookup select Look Up Domain Controllers with DNS, and then select one of the following:
 - User Domain from Login to perform the DNS lookup with the domain name of the login user.
 - o Specify a Domain enter the domain name to use for the DNS lookup
 - To enable CMC to use the specified Active Directory Domain Controller server addresses, select Specify Domain
 Controller Addresses. These are the addresses of the domain controllers where the CMC device object and the
 associated objects are located.
- 4. Click Apply to save the settings.
 - NOTE: You must apply the settings before continuing. If you do not apply the settings, the settings are lost when you navigate to the next page.
- 5. In the Extended Schema Settings section, type the CMC device name and the domain name.
- 6. If you have enabled certificate validation, you must upload the domain forest root certificate authority-signed certificate to CMC. In the **Manage Certificates** section, type the file path of the certificate or browse to the certificate file. Click **Upload** to upload the file to CMC.
 - NOTE: The File Path value displays the relative file path of the certificate you are uploading. You must type the absolute file path, which includes the full path and the complete file name and file extension.

The SSL certificates for the domain controllers must be signed by the root certificate authority-signed certificate. The root certificate authority-signed certificate must be available on the management station accessing CMC.

CAUTION: SSL certificate validation is required by default. Disabling this certificate is not recommended.

7. If you have enabled Single Sign-On (SSO), in the Kerberos Keytab section, click **Browse**, specify the keytab file and click **Upload**. When the upload is complete, a message is displayed indicating a successful or failed upload.

- 8. Click Apply.
 - The CMC web server automatically restarts after you click Apply.
- 9. Log in to the CMC Web interface.
- 10. Select Chassis in the system tree, click the Network tab, and then click the Network subtab. The Network Configuration page is displayed.
- 11. If Use DHCP for CMC Network Interface IP Address is enabled, do one of the following:
 - Select Use DHCP to Obtain DNS Server Addresses to enable the DHCP server to obtain the DNS server addresses automatically.
 - Manually configure a DNS server IP address by leaving the Use DHCP to Obtain DNS Server Addresses check box unchecked and then typing your primary and alternate DNS server IP addresses in the fields provided.
- 12. Click Apply Changes.

The Active Directory settings for extended schema is configured.

Configuring Active Directory With Extended Schema Using RACADM

To configure a CMC Active Directory with Extended Schema by using the RACADM commands, oen a command prompt and enter the following commands at the command prompt:

```
racadm config -g cfgActiveDirectory -o cfgADEnable 1
racadm config -g cfgActiveDirectory -o cfgADType 1
racadm config -g cfgActiveDirectory -o cfgADRacName <RAC common name>
racadm config -g cfgActiveDirectory -o cfgADRacDomain < fully qualified rac domain name >
racadm config -g cfgActiveDirectory -o cfgADDomainController1 < fully qualified domain
name or IP Address of the domain controller >
racadm config -g cfgActiveDirectory -o cfgADDomainController2 < fully qualified domain
name or IP Address of the domain controller >
racadm config -g cfgActiveDirectory -o cfgADDomainController3 < fully qualified domain
name or IP Address of the domain controller >
```

(i) NOTE: You must configure at least one of the three addresses. CMC attempts to connect to each of the configured addresses one-by-one until it makes a successful connection. With Extended Schema, these are the FQDN or IP addresses of the domain controllers where this CMC device is located.

To disable the certificate validation during an handshake (optional):

```
racadm config -g cfgActiveDirectory -o cfgADCertValidationEnable 0
```

(i) NOTE: In this case, you do not have to upload a CA certificate.

To enforce the certificate validation during SSL handshake (optional):

```
racadm config -g cfgActiveDirectory -o cfgADCertValidationEnable 1
```

In this case, you must upload a CA certificate:

```
racadm sslcertupload -t 0x2 -f < ADS root CA certificate >
```

NOTE: If certificate validation is enabled, specify the Domain Controller Server addresses and the FQDN. Make sure that DNS is configured correctly under.

Using the following RACADM command may be optional:

```
racadm sslcertdownload -t 0x1 -f < RAC SSL certificate >
```

Configuring Generic LDAP Users

CMC provides a generic solution to support Lightweight Directory Access Protocol (LDAP)-based authentication. This feature does not require any schema extension on your directory services.

A CMC administrator can now integrate the LDAP server user logins with CMC. This integration requires configuration on both LDAP server and CMC. On the LDAP server, a standard group object is used as a role group. A user who has CMC access becomes a member of the role group. Privileges are still stored on CMC for authorization similar to the working of the Standard Schema setup with Active Directory support.

To enable the LDAP user to access a specific CMC card, the role group name and its domain name must be configured on the specific CMC card. You can configure a maximum of five role groups in each CMC. A user has the option to be added to multiple groups within the directory service. If a user is a member of multiple groups, then the user obtains the privileges of all their groups.

For information about the privileges level of the role groups and the default role group settings, see Types of Users.

Configuring the Generic LDAP Directory to Access CMC

The CMC's Generic LDAP implementation uses two phases in granting access to a user—user authentication, and then the user authorization.

Authentication of LDAP Users

Some directory servers require a bind before a specific LDAP server can be searched for.

To authenticate a user:

- 1. Optionally bind to the Directory Service. The default is an anonymous bind.
 - NOTE: The Windows-based directory servers do not allow anonymous login. Hence, enter the bind DN name and password.
- 2. Search for the user on the basis of the user login. The default attribute is uid. If more than one object is found, then the process returns an error.
- 3. Unbind and perform a bind with the user's DN and password. If the system is unable to bind, then the login will not be successful.
- **4.** If these steps succeed, the user is authenticated.

Authorization Of LDAP Users

To authorize a user:

- 1. Search each configured group for the user's domain name within the member or uniqueMember attributes. An administrator can configure a user domain.
- 2. For every user group that the user belongs to, give the user appropriate user access rights and privileges.

Configuring Generic LDAP Directory Service Using CMC Web Interface

To configure the generic LDAP directory service:

- i NOTE: You must have the Chassis Configuration Administrator privilege.
- 1. In the left pane, click Chassis Overview > User Authentication > Directory Services.
- 2. Select Generic LDAP.
 - The settings to be configured for standard schema is displayed on the same page.
- **3.** Specify the following:
 - i NOTE: For information about the various fields, see the Online Help.
 - Common Settings
 - Server to use with LDAP:
 - o Static server Specify the FQDN or IP address and the LDAP port number.
 - DNS server Specify the DNS server to retrieve a list of LDAP servers by searching for their SRV record within the DNS.

The following DNS query is performed for SRV records:

```
_[Service Name]._tcp.[Search Domain]
```

where <Search Domain> is the root level domain to use within the query and <Service Name> is the service name to use within the query.

For example:

```
_ldap._tcp.dell.com
```

where ldap is the service name and dell.com is the search domain.

- 4. Click **Apply** to save the settings.
 - NOTE: You must apply the settings before continuing. If you do not apply the settings, the settings are lost when you navigate to the next page.
- 5. In the Group Settings section, click a Role Group.
- 6. On the Configure LDAP Role Group page, specify the group domain name and privileges for the role group.
- 7. Click Apply to save the role group settings, click Go Back To Configuration page, and then select Generic LDAP.
- 8. If you have selected Certificate Validation Enabled option, then in the Manage Certificates section, specify the CA certificate to validate the LDAP server certificate during SSL handshake and click Upload. The certificate is uploaded to CMC and the details are displayed.
- Click Apply.
 The generic LDAP directory service is configured.

Configuring Generic LDAP Directory Service Using RACADM

To configure the LDAP directory service, use the objects in cfqLdap and cfqLdapRoleGroup RACADM groups.

There are many options to configure LDAP logins. In most of the cases, some options can be used with their default settings.

NOTE: It is highly recommended to use the racadm testfeature -f LDAP command to test the LDAP settings for first time setups. This feature supports both IPv4 and IPv6.

The required property changes include enabling LDAP logins, setting the server FQDN or IP, and configuring the base DN of the LDAP server.

- \$ racadm config -g cfgLDAP -o cfgLDAPEnable 1
- \$ racadm config -g cfgLDAP -o cfgLDAPServer 192.168.0.1
- \$ racadm config -g cfgLDAP -o cfgLDAPBaseDN dc= company,dc=com

CMC can be configured to optionally query a DNS server for SRV records. If the cfgLDAPSRVLookupEnable property is enabled, the cfgLDAPServer property is ignored. The following query is used to search the DNS for SRV records:

```
_ldap._tcp.domainname.com
```

 $\verb|ldap| in the above query is the \verb|cfgLDAPSRVLookupServiceName| property.$

 $\verb|cfgLDAPSRVLookupDomainName| is configured to be {\bf domainname.com}.$

For more information about the RACADM commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/support/manuals**.

Configuring CMC For Single Sign-On Or Smart Card Login

This section provides information to configure CMC for Smart Card login and Single Sign-On (SSO) login for Active Directory users.

SSO uses Kerberos as an authentication method allowing users, who have signed in as an automatic- or single sign-on to subsequent applications such as Exchange. For single sign-on login, CMC uses the client system's credentials, which are cached by the operating system after you log in using a valid Active Directory account.

Two-factor-authentication, provides a higher-level of security by requiring users to have a password or PIN, and a physical card containing a private key or digital certificate. Kerberos uses this two-factor authentication mechanism allowing systems to prove their authenticity.

NOTE: Selecting a login method does not set policy attributes with respect to other login interfaces, for example, SSH. You must set other policy attributes for other login interfaces also. If you want to disable all other login interfaces, navigate to the **Services** page and disable all (or some) the login interfaces.

Microsoft Windows 2000, Windows XP, Windows Server 2003, Windows Vista, Windows 7, and Windows Server 2008 can use Kerberos as the authentication mechanism for SSO and smart card login.

For information about Kerberos, see the Microsoft Website.

Topics:

- System Requirements
- Prerequisites For Single Sign-On Or Smart Card Login
- Generating Kerberos Keytab File
- Configuring CMC For Active Directory Schema
- Configuring Browser For SSO Login
- Configuring Browser For Smart Card Login
- Configuring CMC SSO Or Smart Card Login For Active Directory Users

System Requirements

To use the Kerberos authentication, the network must include:

- DNS server
- Microsoft Active Directory Server
 - NOTE: If you are using Active Directory on Microsoft Windows 2003, make sure that you have the latest service packs and patched installed on the client system. If you are using Active Directory on Microsoft Windows 2008, make sure that you have installed SP1 along with the following hot fixes:

Windows6.0-KB951191-x86.msu for the KTPASS utility. Without this patch the utility generates bad keytab files.

Windows6.0-KB957072-x86.msu for using GSS_API and SSL transactions during an LDAP bind.

- Kerberos Key Distribution Center (packaged with the Active Directory Server software).
- DHCP server (recommended).
- The DNS server reverse zone must have an entry for the Active Directory server and CMC.

Client Systems

For only Smart Card login, the client system must have the Microsoft Visual C++
 2005 redistributable. For more information see www.microsoft.com/downloads/details.aspx?FamilyID=
 32BC1BEEA3F9-4C13-9C99-220B62A191EE&displaylang=en

• For Single Sign-On or smart card login, the client system must be a part of the Active Directory domain and Kerberos Realm.

CMC

- Each CMC must have an Active Directory account.
- CMC must be a part of the Active Directory domain and Kerberos Realm.

Prerequisites For Single Sign-On Or Smart Card Login

The pre-requisites to configure SSO or Smart Card logins are:

- Set up the kerberos realm and Key Distribution Center (KDC) for Active Directory (ksetup).
- A robust NTP and DNS infrastructure to avoid issues with clock drift and reverse lookup.
- Configure CMC with Active Directory standard schema role group with authorized members.
- For smart card, create Active Directory users for each CMC, configured to use Kerberos DES encryption but not preauthentication.
- Configure the browser for SSO or smart card login.
- Register the CMC users to the Key Distribution Center with Ktpass (this also outputs a key to upload to CMC).

Generating Kerberos Keytab File

To support the SSO and smart card login authentication, CMC supports Windows Kerberos network. The ktpass tool (available from Microsoft as part of the server installation CD/DVD) is used to create the Service Principal Name (SPN) bindings to a user account and export the trust information into a MIT-style Kerberos keytab file. For more information about the ktpass utility, see the Microsoft Website.

Before generating a keytab file, you must create an Active Directory user account for use with the **-mapuser** option of the ktpass command. You must use the same name as the CMC DNS name to which you upload the generated keytab file.

To generate a keytab file using the ktpass tool:

- 1. Run the ktpass utility on the domain controller (Active Directory server), where you want to map CMC to a user account in Active Directory.
- 2. Use the following ktpass command to create the Kerberos keytab file:

```
ktpass -princ HTTP/cmcname.domainname.com@DOMAINNAME.COM -mapuser keytabuser -crypto DES-CBC-MD5 -ptype KRB5_NT_PRINCIPAL -pass * -out c:\krbkeytab
```

NOTE: The cmcname.domainname.com must be in lower case as required by RFC and the @REALM_NAME must be in uppercase. In addition, CMC supports the DES-CBC-MD5 and AES256-SHA1 types of cryptography for Kerberos authentication.

A keytab file is generated that must be uploaded to CMC.

NOTE: The keytab contains an encryption key and must be kept secure. For more information about the *ktpass* utility, see the **Microsoft** Website.

Configuring CMC For Active Directory Schema

For information about configuring CMC for Active Directory standard schema, see Configuring Standard Schema Active Directory.

For information about configuring CMC for Extended Schema Active Directory, see Extended Schema Active Directory Overview.

Configuring Browser For SSO Login

Single Sign-On (SSO) is supported on Internet Explorer versions 6.0 and later, and Firefox versions 3.0 and later.

i NOTE: The following instructions are applicable only if CMC uses Single Sign-On with Kerberos authentication.

Internet Explorer

To configure Internet Explorer for Single Sign-On:

- 1. In the Internet Explorer, select Tools > Internet Options.
- 2. On the Security tab, under Select a zone to view or change security settings, select Local Intranet.
- 3. Click Sites.
 - The Local Intranet dialog box is displayed.
- 4. Click Advanced.
 - The Local Intranet Advance Settings dialog box is displayed.
- 5. In the Add this site to the zone, type the name of CMC and the domain it belongs to and click Add.
 - i NOTE: You can use a wildcard (*) to specify all devices or users in that domain.

Mozilla Firefox

- 1. In Firefox, type about: config in the Address bar.
 - i NOTE: If the browser displays the This might void your warranty warning, click I'll be careful. I promise.
- 2. In the Filter box, type negotiate.
 - The browser displays a list of preference names limited to those containing the word negotiate.
- 3. From the list, double-click network.negotiate-auth.trusted-uris.
- 4. In the Enter string value dialog box, type the CMC's domain name and click OK.

Configuring Browser For Smart Card Login

Internet Explorer — Make sure that the Internet Browser is configured to download Active-X plug-ins.

Configuring CMC SSO Or Smart Card Login For Active Directory Users

You can use CMC web interface or RACADM to configure CMC SSO or smart card login.

Configuring CMC SSO Or Smart Card Login For Active Directory Users Using Web Interface

To configure Active Directory SSO or smart card login for CMC:

- i NOTE: For information about the options, see the Online Help.
- 1. While configuring Active Directory to set up user account, perform the following additional steps:
 - Upload the keytab file.
 - To enable SSO, select the Enable Single Sign-On option.
 - To enable smart card login, select the **Enable Smart-Card Login** option.

NOTE: If these two options are selected, all command line out-of-band interfaces, including secure shell (SSH), Telnet, Serial, and remote RACADM remain unchanged.

2. Click Apply.

The settings are saved.

You can test the Active Directory using Kerberos authentication using the RACADM command:

```
testfeature -f adkrb -u <user>@<domain>
```

where <user> is a valid Active Directory user account.

A command success indicates that CMC is able to acquire Kerberos credentials and access the user's Active Directory account. If the command is not successful, resolve the error and run the command again. For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide on dell.com/support/manuals.

Uploading Keytab File

The Kerberos keytab file serves as the CMC's user name and password credentials to the Kerberos Data Center (KDC), which in turns allows access to the Active Directory. Each CMC in the Kerberos realm must be registered with the Active Directory and must have a unique keytab file.

You can upload a Kerberos Keytab generated on the associated Active Directory Server. You can generate the Kerberos Keytab from the Active Directory Server by executing the ktpass.exe utility. This keytab establishes a trust relationship between the Active Directory Server and CMC.

To upload the keytab file:

- 1. In the left pane, click Chassis Overview > User Authentication > Directory Services.
- 2. Select Microsoft Active Directory (Standard Schema).
- 3. In the Kerberos Keytab section, click Browse, select a keytab file, and click Upload.

When the upload is complete, a message is displayed indicating whether or not the keytab file is successfully uploaded.

Configuring CMC SSO Or Smart Card Login For Active Directory Users Using RACADM

In addition to the steps performed while configuring Active Directory, run the following command to enable SSO:

```
racadm config -g cfgActiveDirectory -o cfgADSSOEnable 1
```

In addition to the steps performed while configuring Active Directory, use the following objects to enable smart card login:

- cfgSmartCardLogonEnable
- cfgSmartCardCRLEnable

Configuring CMC to Use Command Line Consoles

This section provides information about the CMC command line console (or serial/Telnet/Secure Shell console) features, and explains how to set up the system so that you can perform systems management actions through the console. For information about using the RACADM commands in CMC through the command line console, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide*.

Topics:

- CMC Command Line Console Features
- Using Telnet Console With CMC
- Configuring Terminal Emulation Software
- Connecting to Servers or Input Output Module Using Connect Command

CMC Command Line Console Features

The CMC supports the following serial, Telnet, and SSH console features:

- One serial client connection and up to four simultaneous Telnet client connections.
- Up to four simultaneous Secure Shell (SSH) client connections.
- RACADM command support.
- Built-in connect command connecting to the serial console of servers and I/O module; also available as racadm connect.
- Command line editing and history.
- Session timeout control on all console interfaces.

CMC Command Line Interface Commands

When you connect to the CMC command line, you can enter these commands:

Table 30. CMC Command Line Commands

Command	Description
racadm	RACADM commands begin with the keyword racadm, and then followed by a subcommand. For more information, see Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.
connect	Connects to the serial console of a server or I/O module. For more information, see Connecting to Servers or I/O Module Using Connect Command. i NOTE: You can also use the connect RACADM command.
exit, logout, and quit	All the commands perform the same action. They end the current session and return to a login commond line interface.

Using Telnet Console With CMC

You can have up to four Telnet sessions with CMC at a time.

If your management station is running Microsoft Windows XP or Microsoft Windows Server 2003, you may experience an issue with the characters in a CMC Telnet session. This issue may occur as a frozen login where the return key does not respond and the password prompt does not appear.

To fix this issue, download hotfix 824810 from **support.microsoft.com**. For more information, you can also see Microsoft Knowledge Base article 824810.

In the command line interface, you can manage session timeouts using the racadm command, racadm getconfig -g cfgSessionManagement. For more information, see the Chassis Management Controller Version for Dell PowerEdge VRTX Command Line Reference Guide.

Using SSH With CMC

SSH is a command line session that includes the same capabilities as a Telnet session, but with session negotiation and encryption to improve security. CMC supports SSH version 2 with password authentication. SSH is enabled on the CMC by default.

i NOTE: CMC does not support SSH version 1.

When an error occurs during the CMC login, the SSH client issues an error message. The message text is dependent on the client and is not controlled by CMC. Review the RACLog messages to determine the cause of the failure.

NOTE: OpenSSH must be run from a VT100 or ANSI terminal emulator on Windows. You can also run OpenSSH using Putty.exe. Running OpenSSH at the Windows command prompt does not provide full functionality (that is, some keys do not respond and no graphics are displayed). On servers that run Linux, run SSH client services to connect to CMC with any shell.

Four simultaneous SSH sessions are supported at a time. The session timeout is controlled by the cfgSsnMgtSshIdleTimeout property. For more information about the RACADM commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at dell.com/support/Manuals.

CMC also supports Public Key Authentication (PKA) over SSH. This authentication method improves SSH scripting automation by removing the need to embed or prompt for user ID/password. For more information, see Configuring Public Key Authentication over SSH.

SSH is enabled by default. If SSH is disabled, then you can enable it using any other supported interface.

To configure SSH, see Configuring Services.

Supported SSH Cryptography Schemes

To communicate with CMC using SSH protocol, it supports multiple cryptography schemes listed in the following table.

Table 31. Cryptography Schemes

Scheme Type	Scheme
Asymmetric Cryptography	Diffie-Hellman DSA/DSS 512-1024 (random) bits per NIST specification
Symmetric Cryptography	 AES256-CBC RIJNDAEL256-CBC AES192-CBC RIJNDAEL192-CBC AES128-CBC RIJNDAEL128-CBC BLOWFISH-128-CBC 3DES-192-CBC ARCFOUR-128
Message Integrity	 HMAC-SHA1-160 HMAC-SHA1-96 HMAC-MD5-128 HMAC-MD5-96
Authentication	Password

Configure Public Key Authentication Over SSH

You can configure up to six public keys that can be used with the service username over an SSH interface. Before adding or deleting public keys, make sure to use the view command to see what keys are already set up, so that a key is not accidentally overwritten or deleted. The service username is a special user account that can be used when accessing the CMC through SSH. When the PKA over SSH is set up and used correctly, you need not enter username or passwords to log in to the CMC. This can be very useful to set up automated scripts to perform various functions.

(i) NOTE: There is no GUI support for managing this feature, you can use only the RACADM.

When adding new public keys, make sure that the existing keys are not already at the index, where the new key is added. CMC does not perform checks to ensure previous keys are deleted before a new one is added. As soon as a new key is added, it is automatically in effect as long as the SSH interface is enabled.

When using the public key comment section of the public key, remember that only the first 16 characters are utilized by the CMC. The public key comment is used by the CMC to distinguish SSH users when using the RACADM <code>getssninfo</code> command, because all the PKA users use the service username to log in.

For example, if two public keys are set up one with comment PC1 and one with comment PC2:

For more information about the sshpkauth, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Generating Public Keys for Systems Running Windows

Before adding an account, a public key is required from the system that accesses the CMC over SSH. There are two ways to generate the public/private key pair: using PuTTY Key Generator application for clients running Windows or ssh-keygen CLI for clients running Linux.

This section describes simple instructions to generate a public/private key pair for both applications. For additional or advanced usage of these tools, see the application Help.

To use the PuTTY Key Generator to create a basic key for clients running Windows:

- 1. Start the application and select SSH-2 RSA or SSH-2 DSA for the type of key to generate (SSH-1 is not supported).
- 2. Enter the number of bits for the key. RSA key size should be between 768-4096.

(i) NOTE:

- The recommended DSA key length is 1024.
- CMC may not display a message if you add keys less than 768 or greater than 4096, but when you try to log in with these keys, CMC stops responding.
- For DSA keys greater than 2048, use the following RACADM command. CMC accepts RSA keys up to key strength 4096, but the recommended key strength is 1024.

```
racadm -r 192.168.8.14 -u root -p calvin sshpkauth -i svcacct -k 1 -p 0xfff -f dsa_2048.pub
```

3. Click **Generate** and move the mouse in the window as directed.

After the key is created, you can modify the key comment field.

You can also enter a passphrase to make the key secure. Ensure that you save the private key.

- 4. You have two options for using the public key:
 - Save the public key to a file to upload later.
 - Copy and paste the text from the **Public key for pasting** window when adding the account using the text option.

Generating Public Keys for Systems Running Linux

The ssh-keygen application for Linux clients is a command line tool with no graphical user interface. Open a terminal window and at the shell prompt type:

```
ssh-keygen -t rsa -b 1024 -C testing
```

where.

- -t must be dsa or rsa.
- -b specifies the bit encryption size between 768 and 4096.
- -c allows modifying the public key comment and is optional.

The cpassphrase is optional. After the command completes, use the public file to pass to the RACADM for uploading the file.

RACADM Syntax Notes for CMC

When using the racadm sshpkauth command, ensure the following:

- For the -i option, the parameter must be svcacct. All other parameters for -i fail in CMC. The svcacct is a special account for public key authentication over SSH in CMC.
- To log in to the CMC, the user must be service. Users of the other categories do have access to the public keys entered using the sshpkauth command.

Viewing Public Keys

To view the public keys that you have added to the CMC, type:

```
racadm sshpkauth -i svcacct -k all -v
```

To view one key at a time, replace all with a number from 1 - 6. For example, to view key 2, type:

```
racadm sshpkauth -i svcacct -k 2 -v
```

Adding Public Keys

To add a public key to the CMC using the file upload -f option, at the commond line interface console, enter:

```
racadm sshpkauth -i svcacct -k 1 -p 0xfff -f <public key file>
```

NOTE: You can use only the file upload option with remote RACADM. For more information, see Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

To add a public key using the text upload option, enter:

```
racadm sshpkauth -i svcacct -k 1 -p 0xfff -t "<public key text>"
```

Deleting Public Keys

To delete a public key, run the following command:

```
racadm sshpkauth -i svcacct -k 1 -d
```

To delete all public keys, run the following command:

```
racadm sshpkauth -i svcacct -k all -d
```

Configuring Terminal Emulation Software

CMC supports a serial text console from a management station running one of the following types of terminal emulation software:

- Linux Minicom.
- Hilgraeve's HyperTerminal Private Edition (version 6.3).

Complete the tasks in the following subsections to configure the required type of terminal software.

Configuring Linux Minicom

Minicom is a serial port access utility for Linux. The following steps are valid for configuring Minicom version 2.0. Other Minicom versions may differ slightly, but require the same basic settings. To configure other versions of Minicom, see the information in the section Required Minicom Settings of this User's Guide.

Configuring Minicom Version 2.0

NOTE: For best results, set the **cfgSerialConsoleColumns** property to match the number of columns. Be aware that the prompt consumes two characters. For example, for an 80-column terminal window:

```
racadm config -g cfgSerial -o
cfgSerialConsoleColumns 80.
```

- 1. If you do not have a Minicom configuration file, go to the next step. If you have a Minicom configuration file, type minicom<Minicom config file name>, and then go to step 12.
- 2. At the Linux command prompt, type minicom -s.
- 3. Select Serial Port Setup and press <Enter>.
- 4. Press <a>, and then select the appropriate serial device (for example, /dev/ttyS0).
- 5. Press <e>, and then set the Bps/Par/Bits option to 115200 8N1.
- 6. Press <f>, and then set **Hardware Flow Control** to **Yes** and set **Software Flow Control** to **No**. To exit the **Serial Port Setup** menu, press <Enter>.
- 7. Select Modem and Dialing and press <Enter>.
- 8. In the **Modem Dialing and Parameter Setup** menu, press <Backspace> to clear the **init**, **reset**, **connect**, and **hangup** settings so that they are blank, and then press <Enter> to save each blank value.
- 9. When all specified fields are clear, press <Enter> to exit the Modem Dialing and Parameter Setup menu.
- 10. Select Exit From Minicom and press <Enter>.
- 11. At the command shell prompt, type minicom <Minicom config file name>.
- **12.** To exit Minicom, press <Ctrl><a>, <x>, <Enter>.

Make sure that the Minicom window displays a login prompt. When the login prompt appears, your connection is successful. You are now ready to login and access the CMC command line interface.

Required Minicom Settings

See the following table to configure any version of Minicom.

Table 32. Minicom Settings

Setting Description	Required Setting
Bps/Par/Bits	115200 8N1
Hardware flow control	Yes
Software flow control	No
Terminal emulation	ANSI

Table 32. Minicom Settings (continued)

Setting Description	Required Setting
Modem dialing and parameter settings	Clear the init, reset, connect, and hangup settings so that they are blank.

Connecting to Servers or Input Output Module Using Connect Command

CMC can establish a connection to redirect the serial console of a server or I/O module.

For servers, serial console redirection can be accomplished using:

- CMC command line interface (CLI) or the RACADM connect command. For more information about running the RACADM commands, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.
- iDRAC Web interface serial console redirection feature.
- iDRAC Serial Over LAN (SOL) functionality.

In a serial, Telnet, SSH console, CMC supports the connect command to establish a serial connection to a server or I/O module. The server serial console contains both the BIOS boot and setup screens, and the operating system serial console. For the I/O module, the switch serial console is available. There is a single IOM on the chassis.

CAUTION: When run from the CMC serial console, the connect -b option stays connected until the CMC resets. This connection is a potential security risk.

- (i) NOTE: The connect command provides the -b (binary) option. The -b option passes raw binary data, and cfgSerialConsoleQuitKey is not used. Additionally, when connecting to a server using the CMC serial console, transitions in the DTR signal (for example, if the serial cable is removed to connect a debugger) will not result in you exiting the application.
- NOTE: If the IOM does not support console redirection, the connect command displays an empty console. In that case, to return to the CMC console, type the Escape sequence. The default console escape sequence is <Ctrl><\>.

To connect to an IOM:

connect switch-n

where n is an IOM label A1.

When you reference the IOM in the connect command, the IOM is mapped to switch as shown in the following table.

Table 33. Mapping IO Module to Switches

IO Module Label	Switch
A1	switch-a1 or switch- 1

- (i) NOTE: At a time, there can be only one IOM connection per chassis.
- i NOTE: You cannot connect to pass-throughs from the serial console.

To connect to a managed-server serial console, run the command connect server-n, where n is 1-4. You can also use the racadm connect server-n command. When you connect to a server using the -b option, binary communication is assumed and the escape character is disabled. If the iDRAC is not available, the No route to host error message is displayed.

The connect server-n command enables the user to access the server's serial port. After this connection is established, the user can view the server's console redirection through CMC's serial port that includes both the BIOS serial console and the operating system serial console.

NOTE: To view the BIOS boot screens, serial redirection has to be enabled in the servers' BIOS setup. Also, you must set the terminal emulator window to 80×25. Otherwise, the characters on the page are not properly displayed.

NOTE: All keys do not work on the BIOS setup pages. Therefore, provide appropriate keyboard shortcuts for <Ctrl> <Alt> <Delete> and others. The initial redirection screen displays the necessary keyboard shortcuts.

Configuring the Managed Server BIOS for Serial Console Redirection

You can use a Remote Console session to connect to the managed system using the iDRAC web interface (see the iDRAC User's Guide on dell.com/support/manuals).

By default, the Serial communication in the BIOS is turned off. To redirect host text console data to Serial over LAN, you must enable console redirection through COM1. To change the BIOS setting:

- 1. Turn on the managed server.
- 2. Press the <F2> key to enter the BIOS setup utility during POST.
- **3.** Go to **Serial Communication**, and then press <Enter>. In the dialog box, the serial communication list displays the following options:
 - off
 - on without console redirection
 - on with console redirection via COM1

To navigate between these options, press the appropriate arrow keys.

- i NOTE: Make sure that the On with console redirection via COM1 option is selected.
- 4. Enable **Redirection After Boot** (default value is **disabled**). This option enables BIOS console redirection across subsequent reboots.
- 5. Save the changes and exit.

The managed system restarts.

Configuring Windows for Serial Console Redirection

There is no configuration necessary for servers running the Microsoft Windows Server versions, starting with Windows Server 2003. Windows receives information from the BIOS, and enable the Special Administration Console (SAC) console one COM1.

Configuring Linux for Server Serial Console Redirection During Boot

The following steps are specific to the Linux GRand Unified Bootloader (GRUB). Similar changes are necessary for using a different boot loader.

NOTE: When you configure the client VT100 emulation window, set the window or application that is displaying the redirected console to 25 rows×80 columns to make sure proper text is displayed. Else, some text screens will appear distorted.

Edit the /etc/grub.conf file as follows:

1. Locate the general setting sections in the file and type the following two new lines:

```
serial --unit=1 --speed=57600
terminal --timeout=10 serial
```

2. Append two options to the kernel line:

```
kernel console=ttyS1,57600
```

3. If the /etc/grub.conf contains a splashimage directive, comment it out.

The following example shows the changes described in this procedure.

```
# grub.conf generated by anaconda
# Note that you do not have to rerun grub after making
changes
# to this file
# NOTICE: You do not have a /boot partition. This
means that
# all kernel and initrd paths are relative to
/, e.g.
# root (hd0,0)
# kernel /boot/vmlinuz-version ro root=
/dev/sdal
# initrd /boot/initrd-version.img
#boot=/dev/sda
default=0
timeout=10
#splashimage=(hd0,2)/grub/splash.xpm.gz
serial --unit=1 --speed=57600
terminal --timeout=10 serial
title Red Hat Linux Advanced Server (2.4.9-e.3smp)
root (hd0,0)
kernel /boot/vmlinuz-2.4.9-e.3smp ro root=
/dev/sda1 hda=ide-scsi console=ttyS0 console=
ttyS1,57600
initrd /boot/initrd-2.4.9-e.3smp.img
title Red Hat Linux Advanced Server-up (2.4.9-e.3)
root (hd0,00)
kernel /boot/vmlinuz-2.4.9-e.3 ro root=/dev/sda1
initrd /boot/initrd-2.4.9-e.3.img
```

When you edit the /etc/grub.conf file, follow these guidelines:

- Disable GRUB's graphical interface and use the text-based interface. Else, the GRUB screen is not displayed in console redirection. To disable the graphical interface, comment out the line starting with splashimage.
- To start multiple GRUB options to start console sessions through the serial connection, add the following line to all options:

```
console=ttyS1,57600
```

The example shows console=ttyS1,57600 added to only the first option.

Configuring Linux for Server Serial Console Redirection After Boot

Edit the /etc/inittab file as follows:

Add a new line to configure agetty on the COM2 serial port:

```
co:2345:respawn:/sbin/agetty -h -L 57600 ttyS1 ansi
```

The following example shows the file with the new line.

```
#
    inittab This file describes how the INIT process
# should set up the system in a certain
# run-level.
#
    Author: Miquel van Smoorenburg
# Modified for RHS Linux by Marc Ewing and
# Donnie Barnes
#
# Default runlevel. The runlevels used by RHS are:
# 0 - halt (Do NOT set initdefault to this)
# 1 - Single user mode
# 2 - Multiuser, without NFS (The same as 3, if you
```

```
# do not have networking)
# 3 - Full multiuser mode
# 4 - unused
 5 - X11
# 6 - reboot (Do NOT set initdefault to this)
id:3:initdefault:
# System initialization.
si::sysinit:/etc/rc.d/rc.sysinit
10:0:wait:/etc/rc.d/rc 0
11:1:wait:/etc/rc.d/rc 1
12:2:wait:/etc/rc.d/rc 2
13:3:wait:/etc/rc.d/rc 3
14:4:wait:/etc/rc.d/rc 4
15:5:wait:/etc/rc.d/rc 5
16:6:wait:/etc/rc.d/rc 6
# Things to run in every runlevel.
ud::once:/sbin/update
# Trap CTRL-ALT-DELETE
ca::ctrlaltdel:/sbin/shutdown -t3 -r now
# When our UPS tells us power has failed, assume we
have a few
# minutes of power left. Schedule a shutdown for 2
minutes from now.
# This does, of course, assume you have power
installed and your
# UPS is connected and working correctly.
pf::powerfail:/sbin/shutdown -f -h +2 "Power Failure;
System Shutting Down"
# If power was restored before the shutdown kicked in,
cancel it.
pr:12345:powerokwait:/sbin/shutdown -c "Power
Restored; Shutdown Cancelled"
# Run gettys in standard runlevels
co:2345:respawn:/sbin/agetty -h -L 57600 ttyS1 ansi
1:2345:respawn:/sbin/mingetty tty1
2:2345:respawn:/sbin/mingetty tty2
3:2345:respawn:/sbin/mingetty tty3
4:2345:respawn:/sbin/mingetty tty4
5:2345:respawn:/sbin/mingetty tty5
6:2345:respawn:/sbin/mingetty tty6
# Run xdm in runlevel 5
# xdm is now a separate service
x:5:respawn:/etc/X11/prefdm -nodaemon
```

Edit the /etc/securettyfile as follows:

Add a new line, with the name of the serial tty for COM2:

```
ttyS1
```

The following example shows a sample file with the new line.

```
vc/1
vc/2
vc/3
vc/4
vc/5
vc/6
vc/7
vc/8
vc/9
vc/10
vc/11
tty1
tty2
tty3
tty4
tty5
tty6
tty7
```

tty8 tty9 tty10 tty11 ttyS1

Using FlexAddress and FlexAdress Plus

This section provides information about FlexAddress, FlexAddress Plus, and configuring.

i) NOTE: An Enterprise License must be installed on the CMC to use the Flexaddress feature.

Topics:

- About FlexAddress
- Configuring FlexAddress
- Viewing World Wide Name or Media Access Control Addresses
- Viewing WWN or MAC Address Information
- Viewing Basic WWN or MAC Address Information Using Web Interface
- Viewing Advanced WWN or MAC Address Information Using Web Interface
- Viewing WWN or MAC Address Information Using RACADM
- Command Messages
- FlexAddress DELL SOFTWARE LICENSE AGREEMENT

About FlexAddress

If a server is replaced, the FlexAddress for the slot remains the same for the given server slot. If the server is inserted in a new slot or chassis, the server-assigned WWN/MAC is used unless that chassis has the FlexAddress feature enabled for the new slot. If you remove the server, it reverts to the server-assigned address. You need not reconfigure deployment frameworks, DHCP servers, and routers for various fabrics for identifying the new server.

Every server module is assigned unique WWN and/or MAC addresses as part of the manufacturing process. Without FlexAddress, if a server had to be replaced with another server module, the WWN/MAC addresses changes and Ethernet network management tools and SAN resources had to be reconfigured to identify the new server module.

FlexAddress allows CMC to assign WWN/MAC addresses to a particular slot and override the factory addresses. Hence, if the server module is replaced, the slotbased WWN/MACaddresses remain the same. This feature eliminates the need to reconfigure Ethernet network management tools and SAN resources for a new server module.

Also, the *override* action occurs only when a server module is inserted in a FlexAddress enabled chassis; no permanent changes are made to the server module. If a server module is moved to a chassis that does not support FlexAddress, the factory-assigned WWN/MAC addresses are used.

CMC VRTX chassis is shipped with the SD card, which supports FlexAddress, FlexAddress Plus, and Extended Storage features. If the VRTX chassis is shipped with an optional second CMC, the second CMC has an SD card that supports only Extended Storage.

(i) NOTE:

- Data contained on the SD card is encrypted and may not be duplicated or altered in any manner, because it may inhibit system function and cause the system to not function properly.
- The use of an SD card is limited to one chassis only. You cannot use the same SD card on another chassis.

The FlexAddress feature card contains a range of MAC addresses. Before installing FlexAddress, you can determine the range of MAC addresses contained on a FlexAddress feature card by inserting the SD card into a USB Memory Card Reader and viewing the pwwn_mac.xml file. This clear text XML file on the SD card contains an XML tag mac_start that is the first starting hex MAC address that is used for this unique MAC address range. The mac_count tag is the total number of MAC addresses that the SD card allocates. The total MAC range allocated can be determined by:

```
<mac start> + <mac count> - 1 = <mac end>
```

```
(starting_mac)00188BFFDCFA + (mac_count)0xCF - 1 = (ending_mac)00188BFFDDC8
```

NOTE: Lock the SD card prior to inserting in the USB Memory Card Reader to prevent accidentally modifying any contents. Unlock the SD card before inserting into CMC.

About FlexAddress Plus

The FlexAddress Plus is a new feature added to the feature card version 2.0. It is an upgrade from FlexAddress feature card version 1.0. FlexAddress Plus contains more MAC addresses than the FlexAddress feature. Both features allow the chassis to assign World Wide Name/Media Access Control (WWN/MAC) addresses to Fibre Channel and Ethernet devices. Chassis assigned WWN/MAC addresses are globally unique and specific to a server slot.

Viewing FlexAddress Activation Status

A Feature Card contains one or more of the following features: FlexAddress, FlexAddress Plus, and/or Extended Storage.

To view the chassis FlexAddress status using the CMC web interface, click Chassis Overview > Setup.

The General Chassis Settings page is displayed.

The **FlexAddress** has a value **Active** or **Not Active**. The value **Active** indicates that the feature is installed on the chassis and **Not Active** indicates that the feature is not installed and not in use on the chassis.

Run the following RACADM command to view the SD feature card status:

```
racadm featurecard -s
```

The following message is displayed:

```
Active CMC:
The feature card inserted is valid, serial number CNOH871T1374036T00MXA00
The feature card contains the following feature(s)
    FlexAddress: bound
    FlexAddressPlus: bound
    ExtendedStorage: bound
Standby CMC:
The feature card contains the following feature(s)
    ExtendedStorage: bound
```

NOTE: The secondary CMC is optional and the output for the standby CMC is displayed only if the standby CMC is available in the chassis.

Table 34. Status Messages Returned by the featurecard -s Command

Status Message	Actions
No feature card inserted.	Check CMC to verify that the SD card was properly inserted. In a redundant CMC configuration, ensure that the CMC with the SD feature card installed is the active CMC and not the standby CMC.
The feature card inserted is valid and contains the following feature(s) FlexAddress: bound.	No action required.
The feature card inserted is valid and contains the following feature(s) FlexAddress: bound to another chassis, svctag=ABC1234, SD card SN = 1122334455.	Remove the SD card; locate and install the SD card for the current chassis.
The feature card inserted is valid and contains the following feature(s) FlexAddress: not bound.	The feature card can be moved to another chassis or can be reactivated on the current chassis. To reactivate on the current chassis, enter racadm racreset until

Table 34. Status Messages Returned by the featurecard -s Command (continued)

Status Message	Actions
	the CMC module with the feature card installed becomes active.

Use the following RACADM command to display all activated features on the chassis:

```
racadm feature -s
```

The command returns the following status message:

```
Feature Name = FlexAddress

Date/time Activated = 05 Oct 2013 - 11:50:49
Feature installed from SD-card serial number = CN0H871T1374036T00MXA00

Feature Name = FlexAddressPlus
Date/time Activated = 05 Oct 2013 - 11:50:49
Feature installed from SD-card serial number = CN0H871T1374036T00MXA00

Feature Name = ExtendedStorage
Current Status = redundant, active
Date/time Activated = 05 Oct 2013 - 11:50:58
Feature installed from SD-card serial number = CN0H871T1374036T00MXA00
```

If there are no active features on the chassis, the command returns a message:

```
racadm feature -s
No features active on the chassis
```

Dell Feature Cards may contain more than one feature. After any feature included on a Dell Feature Card has been activated on a chassis, any other features that may be included on that Dell Feature Card cannot be activated on a different chassis. In this case, the racadm feature -s command displays the following message for the affected features:

```
ERROR: One or more features on the SD card are active on another chassis
```

For more information about the feature and featurecard commands, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at the support site.

Configuring FlexAddress

FlexAddress is an optional upgrade that allows server modules to replace the factory-assigned WWN/MAC address with a WWN/MAC address provided by the chassis.

- i NOTE: In this section, the term FlexAddress also indicates FlexAddress Plus.
- NOTE: By using the racresetcfg subcommand, you can reset the Flex Address of a CMC to its factory-default setting, which is "disabled". The RACADM syntax is:

```
racadm racresetcfg -c flex
```

For more information about the FlexAddress-related RACADM commands and data about the other factory-default properties, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/cmcmanuals**.

The server must be turned off before you begin configuration. You can enable or disable FlexAddress on a per-fabric-basis. Additionally, you can enable or disable the feature on a per-slot-basis. After you enable the feature on a per-fabric-basis, you can select slots to be enabled. For example, if Fabric-A is enabled, any slots that are enabled have FlexAddress enabled only on Fabric-A. All other fabrics use the factory-assigned WWN/MAC on the server.

NOTE: FlexAddress does not take effect on a server module until the next reboot. When the FlexAddress feature is deployed for the first time on a given server module, it requires a power-down and power-up sequence for FlexAddress to take effect. FlexAddress on Ethernet devices is programmed by the server module BIOS. For the server module BIOS to

program the address, it needs to be operational which requires the server module to be powered up. When the power-down and power-up sequences complete, the chassis-assigned MAC addresses are available for Wake-On-LAN (WOL) function.

Configuring FlexAddress for Chassis-Level Fabric and Slots

At the chassis level, you can enable or disable the FlexAddress feature for fabrics and slots. FlexAddress is enabled on a per-fabric-basis and then slots are selected for participation in the feature. Both fabrics and slots must be enabled to successfully configure FlexAddress.

Configuring FlexAddress for Chassis-Level Fabric and Slots Using CMC Web Interface

If a server is present in the slot, turn it off before enabling the FlexAddress feature on that slot.

To enable or disable fabric and slots to use the FlexAddress feature using the CMC Web interface:

- 1. In the left pane, click Server Overview > Setup > FlexAddress..
- 2. On the **Deploy FlexAddress** page, in the **Select Fabrics for Chassis-Assigned WWN/MACs** section, select the fabric type (**Fabric-A** or **iDRAC**) to which you want to enable FlexAddress. To disable, clear the option.
- 3. On the **Select Slots for Chassis-Assigned WWN/MACs** page, select the **Enabled** option for the slot to which you want to enable FlexAddress. To disable, clear the option.
 - NOTE: Note the following:
 - If a slot is not selected, FlexAddress is not enabled for the selected fabric.
 - When none of the fabrics are selected and a server slot is selected and applied, the following message is displayed
 No fabrics selected! FlexAddress will not be used on this chassis. Select both the fabric
 and the slot to successfully configure FlexAddress.
 - Configuring Flexaddress for slave slot is not allowed. The option is greyed out in the CMC Web interface. The Ethernet devices associated with the slave slot of the server, inherits the master slot configuration.
- 4. To save the settings, click Apply.

Configuring FlexAddress for Chassis-Level Fabric and Slots Using RACADM

To enable or disable fabrics, use the following RACADM command:

```
racadm setflexaddr [-f <fabricName> <state>]
```

where, $\langle fabricName \rangle$ = A or iDRAC and $\langle state \rangle$ = 0 or 1

0 is disable and 1 is enable.

To enable or disable slots, use the following RACADM command:

```
racadm setflexaddr [-i <slot#> <state>
```

where, $\langle slot \# \rangle = 1$ or 4 and $\langle state \rangle = 0$ or 1

0 is disable and 1 is enable.

For more information about the **setflexaddr** command, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

NOTE: If you purchase the FlexAddress or FlexAddressPlus feature with your Dell PowerEdge VRTX, it is preinstalled and enabled for all slots and fabrics. To purchase this feature, contact Dell at dell.com.

NOTE: By using the racresetcfg subcommand, you can reset the Flex Address of a CMC to its factory-default setting, which is "disabled". The RACADM syntax is:

racadm racresetcfg -c flex

For more information about the FlexAddress-related RACADM commands and data about the other factory-default properties, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/cmcmanuals**.

Viewing World Wide Name or Media Access Control Addresses

The **WWN/MAC Summary** page allows you to view the World Wide Name (WWN) configuration and Media Access Control (MAC) address of a slot in the chassis.

Fabric Configuration

The **Fabric Configuration** section displays the type of Input/Output fabric that is installed for Fabric A. A green check mark indicates that the fabric is enabled for FlexAddress. The FlexAddress feature is used to deploy chassis assigned and slot persistent WWN/MAC addresses to various fabrics and slots within the chassis. This feature is enabled on a per-fabric and per-slot—basis.

NOTE: For more information about the FlexAddress feature, see About FlexAddress.

Viewing WWN or MAC Address Information

You can view the WWW/MAC address inventory of Network Adapters for each server slot or all servers in a chassis. The inventory includes the following:

• Fabric Configuration

NOTE:

- Fabric A displays the type of the Input/Output fabric installed. If Fabric A is enabled, unpopulated slots display chassis-assigned MAC addresses for Fabric A.
- o iDRAC management controller is not a fabric, but its FlexAddress is considered as a fabric.
- IF the check box associated with a component is selected, it implies that the fabric is enabled for FlexAddress or FlexAddressPlus.
- Protocol that is being used on the NIC Adapter port. For example, LAN, ISCI, and FCoE.
- Fibre Channel World Wide Name (WWN) configuration and Media Access Control (MAC) addresses of a slot in the chassis.
- MAC address assignment type and the current active address type Server assigned, FlexAddress, or I/O Identity MAC. A black check mark indicates the active address type, either server-assigned, chassis-assigned, or remote assigned.
- Status of the NIC partitions for devices supporting partitioning.

You can view the WWN/MAC Address inventory using the Web interface or the RACADM CLI. Based on the interface, you can filter the MAC address and know which WWN/MAC address is in use for that function or partition. If the adapter has NPAR enabled, you can view which partitions are enabled or disabled.

Using the Web interface, you can view the WWN/MAC Addresses information for:

- Specific slots Open the FlexAddress page, by clicking Server Overview > Slot <x> > Setup > FlexAddress.
- All the slots and server Open the WWN/MAC Summary page, by clicking Server Overview > Properties > WWN/

From both the pages, you can view the WWN/MAC Addresses information in the basic mode or the advanced mode:

- Basic Mode In this mode, you can view Server Slot, Fabric, Protocol, WWN/MAC addresses, and Partition Status. Only Active MAC addresses are displayed in the WWN/MAC address box. You can filter using any or all of the fields displayed:
- Advanced Mode In this mode you can view all the fields displayed in the basic mode and all the MAC types (Server Assigned, Flex Address, and IO Identity). You can filter using any or all of the fields displayed.

In both the Basic mode and the Advanced mode, the WWN/MAC Addresses information is displayed in a collapsed form. Click the plus symbol corresponding to a slot or click **Expand/Collapse All** to view the information for a specific slot or all the slots.

You can also export the WWN/MAC Addresses information for all the servers in the chassis to a local folder.

For information about the fields, see the Online Help.

Viewing Basic WWN or MAC Address Information Using Web Interface

To view WWN/MAC Address information for each server slot or all servers in a chassis, in the basic mode:

- Click Server Overview > Properties > WWN/MAC
 The WWN/MAC Summary page displays the WWN/MAC Address Information.
 - Alternatively, click **Server Overview** > **Slot** <**x>** > **Setup** > **FlexAddress** to view the WWN/MAC Address information for a specific server slot. The **FlexAddress** page is displayed.
- 2. In the WWN/MAC Addresses table, click Export to save the WWN/MAC addresses locally.
- 3. Click the plus symbol corresponding to a slot or click **Expand/Collapse All** to expand or collapse the attributes listed for a specific slot or all the slots in the WWN/MAC Addresses table.
- 4. From the View drop-down menu, select Basic, to view the WWN/MAC Addresses attributes in tree view.
- **5.** From the **Server Slot** drop-down menu, select **All Servers** or a specific Slot to view the WWN/MAC Addresses attributes for all servers or servers in specific slots only respectively.
- 6. From the **Fabric** drop-down menu, select one of the fabric types to view details for all or specific type of management or I/O fabric associated with the servers.
- 7. From the **Protocol** drop-down menu, select **All Protocols** or one of the listed network protocols to view all the MACS or the MACs associated with the selected protocol.
- 8. In the **WWN/MAC Addresses** field, enter a partial MAC address or the full MAC address to view only the slots associated with the specific MAC address.
- 9. From the **Partition Status** drop-down menu, select the status of the partitions to display servers with the selected partition status.

For information about the fields, see the Online Help.

Viewing Advanced WWN or MAC Address Information Using Web Interface

To view WWN/MAC Address Information for each server slot or all servers in a chassis, in the advanced mode:

- Click Server Overview > Properties > WWN/MAC
 The WWN/MAC Summary page displays the WWN/MAC Address Information.
- 2. From the View drop-down menu, select Advanced, to view the WWN/MAC Addresses attributes in detailed view. In the WWN/MAC Addresses table displays Server Slot, Fabric, Protocol, WWN/MAC addresses, Partition Status, and the MAC address assignment type Server assigned, FlexAddress, or I/O Identity MAC. A black check mark indicates the active address type, either server-assigned, chassis-assigned, or remote assigned. MAC.
- 3. In the WWN/MAC Addresses table, click Export to save the WWN/MAC addresses locally.
- 4. Click the against a slot or click **Expand/Collapse All** to expand or collapse the attributes listed for a specific slot or all the slots in the WWN/MAC Addresses table.
- 5. From the **Server Slot** drop-down menu, select **All Servers** or a specific Slot to view the WWN/MAC Addresses attributes for all servers or servers in specific slots only respectively.

- 6. From the **Fabric** drop-down menu, select one of the fabric types to view details for all or specific type of management or I/O fabric associated with the servers.
- 7. From the **Protocol** drop-down menu, select **All Protocols** or one of the listed network protocols to view all the MACS or the MACs associated with the selected protocol.
- 8. In the **WWN/MAC Addresses** field, enter the MAC address to view only the slots associated with the specific MAC address.
- 9. From the **Partition Status** drop-down menu, select the status of the partitions to display servers with the selected partition status.

If a particular partition is disabled, the status is displayed as **Disabled** and the row displaying the partition is greyed out.

For information about the fields, see the Online Help.

Viewing WWN or MAC Address Information Using RACADM

To view WWN/MAC address information for all servers or specific servers using RACADM, use the getflexaddr and getmacaddress subcommands.

To display Flexaddress for the entire chassis, use the following RACADM command:

```
racadm getflexaddr
```

To display Flexaddress status for a particular slot, use the following RACADM command:

```
racadm getflexaddr [-i <slot#>]
```

where <slot #> is a value from 1 to 4.

To display the NDC or LOM MAC address, use the following RACADM command:

```
racadm getmacaddress
```

To display the MAC address for chassis, use the following RACADM command:

```
racadm getmacaddress -m chassis
```

To display the iSCSI MAC addresses for all servers, use the following RACADM command:

```
racadm getmacaddress -t iscsi
```

To display the iSCSI MAC for a specific server, use the following RACADM command:

```
racadm getmacaddress [-m <module> [-x]] [-t iscsi]
```

To display the user-defined MAC and WWN address, use the following RACADM command:

```
racadm getmacaddress -c io-identity
```

```
racadm getmacaddress -c io-identity -m server -2
```

To display the console assigned MAC/WWN of all LOMs or mezzanine cards, use the following RACADM command:

```
racadm getmacaddress -c all
```

To display the chassis assigned WWN/MAC address, use the following RACADM command:

```
racadm getmacaddress -c flexaddress
```

To display the MAC/WWN addresses for all LOMs or mezzanine cards, use the following RACADM command:

```
racadm getmacaddress -c factory
```

To display the Ethernet and iSCSI MAC/WWN addresses for all iDRAC/LOMs/mezzanine cards, use the following RACADM command:

racadm getmacaddress -a

For more information on the getflexaddr and getmacaddress subcommand, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Command Messages

The following table lists the RACADM commands and output for common FlexAddress situations.

Table 35. FlexAddress Commands and Output

Situation	Command	Output
SD card in the active CMC module is bound to another service tag.	\$racadm featurecard -s	The feature card inserted is valid and contains the following feature(s)
		FlexAddress: bound to another chassis, svctag = <service number="" tag=""> SD card SN = <valid address="" flex="" number="" serial=""></valid></service>
SD card in the active CMC module that is bound to the same service tag.	\$racadm featurecard -s	The feature card inserted is valid and contains the following feature(s)
		FlexAddress: bound
SD card in the active CMC module that is not bound to any service tag.	\$racadm featurecard -s	The feature card inserted is valid and contains the following feature(s)
		FlexAddress:not bound
FlexAddress feature not active on the chassis for any reason (No SD card inserted/ corrupt SD card/ after feature deactivated /SD card bound to a different chassis).	<pre>\$racadm setflexaddr [-f <fabricname> <slotstate>] \$racadm setflexaddr [-i <slot#> <slotstate>]</slotstate></slot#></slotstate></fabricname></pre>	ERROR: Flexaddress feature is not active on the chassis
Guest user attempts to set FlexAddress on slots/fabrics.	<pre>\$racadm setflexaddr [-f <fabricname> <slotstate>] \$racadm setflexaddr [-i <slot#> <slotstate>]</slotstate></slot#></slotstate></fabricname></pre>	ERROR: Insufficient user privileges to perform operation
Deactivating FlexAddress feature with chassis powered ON.	\$racadm feature -d -c flexaddress	ERROR: Unable to deactivate the feature because the chassis is powered ON
Guest user tries to deactivate the feature on the chassis.	\$racadm feature -d -c flexaddress	ERROR: Insufficient user privileges to perform operation
Changing the slot/fabric FlexAddress settings while the server modules are powered ON.	\$racadm setflexaddr -i 1 1	ERROR: Unable to perform the set operation because it affects a powered ON server

Table 35. FlexAddress Commands and Output (continued)

Situation	Command	Output
Changing the Flexaddress settings of slot or fabric, when the CMC Enterprise License is not installed.	<pre>\$racadm setflexaddr - i<slotnum> <status></status></slotnum></pre>	ERROR: SWC0242 : A required license is missing or expired.
	<pre>\$racadm setflexaddr - f<fabricname> <status></status></fabricname></pre>	Obtain an appropriate license and try again, or contact your service provider for additional details.
		NOTE: To resolve this issue, you must have a FlexAddress Enablement license.

FlexAddress DELL SOFTWARE LICENSE AGREEMENT

This is a legal agreement between you, the user, and Dell Products L.P. or Dell Global B.V. ("Dell"). This agreement covers all software that is distributed with the Dell product, for which there is no separate license agreement between you and the manufacturer or owner of the software (collectively the "Software"). This agreement is not for the sale of Software or any other intellectual property. All title and intellectual property rights in and to Software is owned by the manufacturer or owner of the Software. All rights not expressly granted under this agreement are reserved by the manufacturer or owner of the Software. By opening or breaking the seal on the Software packet(s), installing or downloading the Software, or using the Software that has been preloaded or is embedded in your product, you agree to be bound by the terms of this agreement. If you do not agree to these terms, promptly return all Software items (disks, written materials, and packaging) and delete any preloaded or embedded Software.

You may use one copy of the Software on only one computer at a time. If you have multiple licenses for the Software, you may use as many copies at any time as you have licenses. "Use" means loading the Software in temporary memory or permanent storage on the computer. Installation on a network server solely for distribution to other computers is not "use" if (but only if) you have a separate license for each computer to which the Software is distributed. You must ensure that the number of persons using the Software installed on a network server does not exceed the number of licenses that you have. If the number of users of Software installed on a network server exceeds the number of licenses, you must purchase additional licenses until the number of licenses equals the number of users before allowing additional users to use the Software. If you are a commercial customer of Dell or a Dell affiliate, you hereby grant Dell, or an agent selected by Dell, the right to perform an audit of your use of the Software during normal business hours, you agree to cooperate with Dell in such audit, and you agree to provide Dell with all records reasonably related to your use of the Software. The audit is limited to verification of your compliance with the terms of this agreement.

The Software is protected by United States copyright laws and international treaties. You may make one copy of the Software solely for backup or archival purposes or transfer it to a single hard disk provided you keep the original solely for backup or archival purposes. You may not rent or lease the Software 240 Using FlexAddress and FlexAdress Plus Cards or copy the written materials accompanying the Software, but you may transfer the Software and all accompanying materials on a permanent basis as part of a sale or transfer of the Dell product if you retain no copies and the recipient agrees to the terms hereof. Any transfer must include the most recent update and all prior versions. You may not reverse engineer, decompile or disassemble the Software. If the package accompanying your computer contains compact discs, 3.5" and/or 5.25" disks, you may use only the disks appropriate for your computer. You may not use the disks on another computer or network, or loan, rent, lease, or transfer them to another user except as permitted by this agreement.

LIMITED WARRANTY

Dell warrants that the Software disks is free from defects in materials and workmanship under normal use for ninety (90) days from the date you receive them. This warranty is limited to you and is not transferable. Any implied warranties are limited to ninety (90) days from the date you receive the Software. Some jurisdictions do not allow limits on the duration of an implied warranty, so this limitation may not apply to you. The entire liability of Dell and its suppliers, and your exclusive remedy, shall be (a) return of the price paid for the Software or (b) replacement of any disk not meeting this warranty that is sent with a return authorization number to Dell, at your cost and risk. This limited warranty is void if any disk damage has resulted from accident, abuse, misapplication, or service or modification by someone other than Dell. Any replacement disk is warranted for the remaining original warranty period or thirty (30) days, whichever is longer.

Dell does NOT warrant that the functions of the Software meets your requirements or that operation of the Software is uninterrupted or error free. You assume responsibility for selecting the Software to achieve your intended results and for the use and results obtained from the Software.

DELL, ON BEHALF OF ITSELF AND ITS SUPPLIERS, DISCLAIMS ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR

PURPOSE, FOR THE SOFTWARE AND ALL ACCOMPANYING WRITTEN MATERIALS. This limited warranty gives you specific legal rights; you may have others, which vary from jurisdiction to jurisdiction.

IN NO EVENT SHALL DELL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR OTHER PECUNIARY LOSS) ARISING OUT OF USE OR INABILITY TO USE THE SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Because some jurisdictions do not allow an exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

OPEN SOURCE SOFTWARE

A portion of this CD may contain open source software, which you can use under the terms and conditions of the specific license under which the open source software is distributed.

THIS OPEN SOURCE SOFTWARE IS DISTRIBUTED IN THE HOPE THAT IT WILL BE USEFUL, BUT IS PROVIDED "AS IS" WITHOUT ANY EXPRESSED OR IMPLIED WARRANTY; INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL DELL, THE COPYRIGHT HOLDERS, OR THE CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTUTUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILTLY, WHETHER IN CONTRACT, STRICT LIABITLY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILTLY OF SUCH DAMAGE.

U.S. GOVERNMENT RESTRICTED RIGHTS

The software and documentation are "commercial items" as that term is defined at 48 C.F.R. 2.101, consisting of "commercial computer software" and "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212. Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4, all U.S. Government end users acquire the software and documentation with only those rights set forth herein.

Contractor/manufacturer is Dell Products, L.P., One Dell Way, Round Rock, Texas 78682.

GENERAL

This license is effective until terminated. It terminates upon the conditions set forth above or if you fail to comply with any of its terms. Upon termination, you agree that the Software and accompanying materials, and all copies thereof, is destroyed. This agreement is governed by the laws of the State of Texas. Each provision of this agreement is severable. If a provision is found to be unenforceable, this finding does not affect the enforceability of the remaining provisions, terms, or conditions of this agreement. This agreement is binding on successors and assigns. Dell agrees and you agree to waive, to the maximum extent permitted by law, any right to a jury trial with respect to the Software or this agreement. Because this waiver may not be effective in some jurisdictions, this waiver may not apply to you. You acknowledge that you have read this agreement, that you understand it, that you agree to be bound by its terms, and that this is the complete and exclusive statement of the agreement between you and Dell regarding the Software.

Managing Fabrics

The chassis supports a fabric type, which is Fabric A. Fabric A is used by the single I/O Module, and is always connected to the on-board Ethernet adapters of the servers.

The chassis has only one I/O module (IOM), where the IOM is a pass-through or switch module. The I/O Module is classified as group A.

Chassis IOM uses a discrete data path called **Fabric**, and it is named A. The Fabric A supports only Ethernet. Each server IO adapter (Mezzanine Card or LOM) can have either two or four ports depending on the capability. The mezzanine card slots are occupied by PCIe extension cards that are connected to PCIe cards (and not to IO modules). When you deploy the Ethernet, iSCSI, or FibreChannel networks, span their redundant links across banks one and two for maximum availability. The discrete IOM is identified with a fabric identifier.

NOTE: In the CMC CLI, the IOM is referred to by the convention, switch.

Topics:

- Fresh Power-up Scenario
- Monitoring IOM Health
- Configuring Network Settings for IOM
- Managing Power Control Operation for IOMs
- Enabling or Disabling LED Blinking for IOMs

Fresh Power-up Scenario

When the chassis is plugged in and turned on, the I/O module has priority over the servers. The IOM is allowed to turn on before the others. At this time, verification of their fabric types is not performed.

After the IOMs turn on, the servers turn on, and then CMC verifies the servers for fabric consistency.

A pass-through module and switch are allowed in the same group if their fabric is identical. Switches and pass-through modules can exist in the same group even if they are manufactured by different vendors.

Monitoring IOM Health

For information about monitoring IOM health, see Viewing Information and Health Status of the IOM.

Configuring Network Settings for IOM

You can specify the network settings for the interface used to manage the IOM. For Ethernet switches, the out-of-band management port (IP address) is configured. The in-band management port (that is, VLAN1) is not configured using this interface.

Before configuring the network settings for the IOM, make sure the IOM is turned on.

To configure the network setting of IOM in Group A, you must have the Fabric A Administrator privileges.

- NOTE: For Ethernet switches, the in-band (VLAN1) and out-of-band management IP addresses cannot be the same, or cannot be on the same network. This results in the out-of-band IP address in being not set. See the IOM documentation for the default in-band management IP address.
- (i) NOTE: Do not configure I/O module network settings for Ethernet pass-through and Infiniband switches.

Configuring Network Settings for IOM Using CMC Web Interface

To configure the network settings for I/O Module:

- In the left pane, click Chassis Overview, click I/O Module Overview, and then click Setup. Alternatively, to configure the network settings of the only available I/O module that is A, click A Gigabit Ethernet, and then click Setup.
 On the Configure I/O Module Network Settings page, type appropriate data, and then click Apply.
- 2. If allowed, type the root password, SNMP RO Community string, and Syslog Server IP Address for the IOM. For more information about the field descriptions, see the *Online Help*.
 - NOTE: The IP address set on the IOM from CMC is not saved to the permanent startup configuration of the switch. To permanently save the IP address configuration, you must run the connect switch command, or racadm connect switch RACADM command, or use a direct interface to the IOM GUI to save this address to the startup configuration file.
 - (i) NOTE: The length of the SNMP community string can be in the ASCII value range of 33–125 characters.
- 3. Click Apply.

The network settings are configured for the IOM.

(i) NOTE: If allowed, you can reset the VLANs, network properties, and IO ports to its default configuration values.

Configuring Network Settings for IOM Using RACADM

To configure the network settings for an IOM by using RACADM, set the date and time. See the deploy command section in the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

You can set the user name, password, and SNMP string for the IOM using the RACADM deploy command:

```
racadm deploy -m switch -u <username> -p <password>
racadm deploy -m switch -u -p <password> -v SNMPv2 <snmpCommunityString> ro
racadm deploy -a [server|switch] -u <username> -p <password>
```

Managing Power Control Operation for IOMs

For information to set the power control operation for IOMs, see Executing Power Control Operations on the IOM.

Enabling or Disabling LED Blinking for IOMs

For information to enable LED blinking for IOMs, see Configuring LEDs to Identify Components on the Chassis.

Managing and Monitoring Power

The PowerEdge VRTX chassis is the most power-efficient modular server enclosure. It is designed to include highly efficient power supplies and fans, has an optimized layout for the air to flow more easily through the system, and contains power-optimized components throughout the enclosure. The optimized hardware design is coupled with sophisticated power management capabilities that are built into the Chassis Management Controller (CMC), power supplies, and iDRAC to allow you to further enhance power-efficient server environment.

The Power Management features of the PowerEdge VRTX help administrators configure the enclosure to reduce power consumption and to adjust the power as required specific to the environment.

The PowerEdge VRTX modular enclosure consumes AC power and distributes the load across all active internal power supply units (PSUs). The system can deliver up to 4800 Watts of AC power that is allocated to server modules and the associated enclosure infrastructure. However, this capacity varies based on the power redundancy policy that you select.

The PowerEdge VRTX enclosure can be configured for any of the two redundancy policies that affect PSU behavior and determine how chassis Redundancy state is reported to administrators.

You can also control Power management through **OpenManage Power Center (OMPC)**. When OMPC controls power externally, CMC continues to maintain:

- Redundancy policy
- Remote power logging
- Dynamic power supply engagement (DPSE)

OMPC then manages:

- Server power
- Server priority
- System Input Power Capacity
- Maximum Power Conservation Mode

i NOTE: Actual power delivery is based on configuration and workload.

You can use the CMC web interface or RACADM to manage and configure power controls on CMC:

- View power allocations, consumption, and status for the chassis, servers, and PSUs.
- Configure power budget and redundancy policy for the chassis.
- Execute power control operations (turn on, turn off, system reset, power-cycle) for the chassis.

Topics:

- Redundancy Policies
- Dynamic Power Supply Engagement
- Default Redundancy Configuration
- Power Budgeting For Hardware Modules
- Server Slot Power Priority Settings
- Assigning Priority Levels To Servers
- Assigning Priority Levels To Servers Using CMC Web Interface
- Assigning Priority Levels To Servers Using RACADM
- Viewing Power Consumption Status
- Viewing Power Budget Status Using CMC Web Interface
- Redundancy Status and Overall Power Health
- Configuring power budget and redundancy
- Executing Power Control Operations
- Executing Power Control Operations on a Server
- Executing Power Control Operations for Multiple Servers Using CMC Web Interface
- Executing Power Control Operations on the IOM

Redundancy Policies

Redundancy policy is a configurable set of properties that determine how CMC manages power to the chassis. The following redundancy policies are configurable with or without dynamic PSU engagement:

- Grid redundancy
- Power supply redundancy

Grid Redundancy Policy

The purpose of the Grid redundancy policy is to enable a modular enclosure system to operate in a mode in which it can tolerate AC power failures. These failures may originate in the AC power grid, the cabling and delivery, or a PSU itself.

When you configure a system for Grid redundancy, the PSUs are divided into grids: PSUs in slots 1 and 2 are in the first grid while PSUs in slots 3 and 4 are in the second grid. CMC manages power so that if there is a failure of either grid, the system continues to operate without any degradation. Grid redundancy also tolerates failures of individual PSUs.

- NOTE: One role of Grid redundancy is to provide seamless server operation despite failure of a whole power grid, but the most power is available to maintain Grid redundancy when the capacities of the two grids are approximately equal.
- i NOTE: Grid redundancy is only met when the load requirements do not exceed the capacity of the weakest power grid.

Grid Redundancy Levels

One PSU in each grid is the minimum configuration necessary for use as grid redundant. Additional configurations are possible with every combination that has at least one PSU in each grid. However, to make the maximum power available for use, the total power of the PSUs in each leg should be as close to equal as practical. The upper limit of power while maintaining Grid redundancy is the power available on the weakest of the two grids.

If a CMC is unable to maintain Grid redundancy, an e-mail and/or SNMP alert is sent to administrators, if the Redundancy Lost event is configured for alerting.

In the event of a single PSU not functioning in this configuration, the remaining PSUs in the problematic grid are marked as online. In this state, the PSUs in the Redundant Grid if not in failed state, help in functioning of the system without interruption. If a PSU stops functioning, the chassis health is marked non-critical. If the smaller grid cannot support the total chassis power allocations, then Grid redundancy status is reported as **No**, and the health of chassis is displayed as **Critical**.

Power Supply Redundancy Policy

The power supply redundancy policy is useful when redundant power grids are not available, but you may want to be protected against a single PSU failure bringing down your servers in a modular enclosure. The highest capacity PSU is kept in online reserve for this purpose. This forms a Power Supply redundancy pool.

PSUs beyond those required for power and redundancy are still available and is added to the pool in the event of a failure.

Unlike Grid redundancy, when power supply redundancy is selected, CMC does not require the PSU units to be present in any specific PSU slot positions.

NOTE: Dynamic Power Supply Engagement (DPSE) allows PSUs to be placed in standby. The standby state indicates a physical state of PSUs that are not supplying power. When you enable DPSE, the extra PSUs may be placed in Standby mode to increase efficiency and save power.

Dynamic Power Supply Engagement

By default, Dynamic Power Supply Engagement (DPSE) mode is disabled. DPSE saves power by optimizing the power efficiency of the PSUs that supply power to the chassis. This also results in increased PSU life, and reduced heat generation. To use this feature, you must have an Enterprise License.

CMC monitors total enclosure power allocation, and moves the PSUs into Standby state, causing the total power allocation of the chassis to be delivered through fewer PSUs. As the online PSUs are more efficient when running at higher utilization, it improves their efficiency and longevity of the standby PSUs.

To operate remaining PSUs at their maximum efficiency, use the following power redundancy modes:

- **PSU Redundancy** mode with DPSE provides power efficiency. At least two supplies are online, with one PSU required to power the configuration, and one to provide redundancy in case of a PSU failure. PSU Redundancy mode offers protection against the failure of any one PSU, but offers no protection in the event of an AC grid loss.
- **Grid Redundancy** mode with DPSE, where at least two PSUs are active, one on each power grid. Grid redundancy also balances the efficiency and maximum availability for a partially-loaded modular enclosure configuration.
- Disabling DPSE provides the lowest efficiency as all four supplies are active and share the load, resulting in lower utilization of each power supply.

DPSE can be enabled for all two power supply redundancy configurations explained above — **Power Supply Redundancy**, and **Grid Redundancy**.

- i) NOTE: In a two PSU configuration modes, server load may prevent any PSU from changing to the Standby mode.
- In a **Power Supply Redundancy** configuration, in addition to the PSUs required to power the enclosure, the enclosure always keeps an additional PSU powered on and marked **Online**. Power utilization is monitored and one PSU can be moved to Standby state on the basis of overall system load. In a four PSU configuration, a minimum of two PSUs are always turned on.

Because an enclosure in the **Power Supply Redundancy** configuration always has an extra PSU engaged, the enclosure can accommodate the loss of one online PSU and still have enough power for the installed server modules. The loss of the online PSU causes a standby PSU to come online. Simultaneous failure of multiple PSUs may result in the loss of power to some server modules while the standby PSUs are turning on.

• In **Grid Redundancy** configuration, all PSUs are engaged when the chassis is turned on. Power utilization is monitored, and if system configuration and power utilization allows, PSUs are moved to the **Standby** state. Because the **Online** status of PSUs in a grid mirrors that of the other grid, the enclosure can sustain the loss of power to an entire grid with no interruption of power to the enclosure.

An increase in power demand in the **Grid Redundancy** configuration cause the engagement of PSUs from the **Standby** state. This maintains the mirrored configuration needed for dual-grid redundancy.

NOTE: With DPSE in enabled state, if power demand increases in both the two Power Redundancy policy modes, the standby PSUs are brought **Online** to reclaim power.

Default Redundancy Configuration

As shown in the table here, the default redundancy configuration for a chassis depends on the number of PSUs that it contains.

Table 36. Default Redundancy Configuration

PSU Configuration	Default Redundancy Policy	Default Dynamic PSU Engagement Setting
Two PSUs	DC Redundancy	Disabled
Four PSUs	DC Redundancy	Disabled

Grid Redundancy

In Grid redundancy mode with four PSUs, all four PSUs are active. The two PSUs must connect to one AC power grid, while the other two PSUs are connect to the other AC power grid.

CAUTION: To avoid a system failure, and for Grid redundancy to work effectively, there must be a balanced set of PSUs properly cabled to separate AC grids.

If one AC grid fails, the PSUs on the functioning AC grid take over, without interrupting the servers or infrastructure.

CAUTION: In Grid redundancy mode, you must have balanced sets of PSUs (at least one PSU in each grid). If this condition is not met, Grid redundancy is not possible.

Power Supply Redundancy

When power supply redundancy is enabled, a PSU in the chassis is kept as a spare, ensuring that the failure of any one PSU does not cause the servers or chassis to turn off. Power supply redundancy mode requires minimum two PSUs. Additional PSUs, if present, are utilized to improve power efficiency of the system if DPSE is enabled. Subsequent failures after loss of redundancy may cause the servers in the chassis to turn off.

Power Budgeting For Hardware Modules

CMC offers a power budgeting service that allows you to configure power budget, redundancy, and dynamic power for the chassis.

The power management service enables optimization of power consumption and reallocation of power to different modules on the basis of demand.

CMC maintains a power budget for the enclosure that reserves the necessary wattage for all installed servers and components.

CMC allocates power to the CMC infrastructure and the servers in the chassis. CMC infrastructure consists of components in the chassis, such as fans, I/O module, and storage adapters, PCle cards, physical disk, main board. The chassis may have up to four servers that communicate to the chassis through an iDRAC. For more information, see the *iDRAC User's Guide* at **dell.com/support/manuals**.

iDRAC provides CMC with its power envelope requirements before powering up the server. The power envelope consists of the maximum and minimum power requirements necessary to keep the server operating. iDRAC's initial estimate is based on its initial understanding of components in the server. After operation commences and further components are discovered, iDRAC may increase or decrease its initial power requirements.

When a server is turned on in an enclosure, the iDRAC software reestimates the power requirements and requests a subsequent change in the power envelope.

CMC supplies the requested power to the server, and the allocated wattage is subtracted from the available budget. After the server is granted a power request, the server's iDRAC software continuously monitors the actual power consumption. On the basis of actual power requirements, the iDRAC power envelope may change over a period of time. iDRAC requests a power step up if the servers are fully using the allocated power.

Under heavy load, the performance of the processors on the server may be degraded to ensure power consumption stays lower than the user-configured System Input Power Cap.

The PowerEdge VRTX enclosure can supply enough power for peak performance of most server configurations, but many available server configurations do not consume the maximum power that the enclosure can supply. To help datacenters allocate power for their enclosures, the PowerEdge VRTX allows you to specify a System Input Power Cap to make sure that the overall chassis AC power draw stays within a given threshold point. CMC first makes sure that enough power is available to run the fans, I/O module, storage adapters, physical disk drive, main board, and CMC itself. This power allocation is called the Input Power Allocated to Chassis Infrastructure. After Chassis infrastructure, the servers in an enclosure are turned on. Any attempt to set a System Input Power Cap less than the "Power Burden" will not be successful. Power Burden is the sum of power allocated to the infrastructure and the minimum power allocated for the powered servers.

(i) NOTE: To use the Power Cap feature, you must have an Enterprise License.

If necessary for the total power budget to stay below the value of the *System Input Power Cap*, CMC allocates servers a value less than their maximum requested power. Servers are allocated power based on their *Server Priority* setting, with higher priority servers getting maximum power, priority 2 servers getting power after priority 1 servers, and so on. Lower priority servers may get less power than priority-one servers based on *System Input Max Power Capacity*, and the user-configured setting of *System Input Power Cap*.

Configuration changes, such as an additional server, shared HDDs, or PCle cards in the chassis, may require the *System Input Power Cap* to be increased. Power needs in a modular enclosure also increase when thermal conditions change and the fans are required to run at higher speed, which causes them to consume additional power. Insertion of I/O module and storage adapters, PCle cards, physical disk, main board; number, type, and configuration of PSUs also increase the power needs of the modular enclosure. A fairly small amount of power is consumed by servers even when they are powered down to keep the management controller powered up.

Additional servers can be powered up in the modular enclosure only if sufficient power is available. The *System Input Power Cap* can be increased any time up to a maximum value of 5000 Watt to allow the power up of additional servers.

Changes in the modular enclosure that reduce the power allocation are:

• Server turned off

- I/O module turned off
- Storage adapters, PCle cards, physical disk drive, and main board turned off
- Transition of the chassis to a turned-off state

You can reconfigure the System Input Power Cap when the chassis is either turned on or turned off.

Server Slot Power Priority Settings

CMC allows you to set a power priority for each of the four server slots in an enclosure. The priority settings are 1 (highest) through 9 (lowest). These settings are assigned to slots in the chassis, and the priority of the slots is inherited by any server inserted in that slot. CMC uses slot priority to preferentially budget power to the highest priority servers in the enclosure.

According to the default server slot priority setting, power is equally apportioned to all slots. Changing the slot priorities allows administrators to prioritize the servers that are given preference for power allocations. If the more critical server modules are left at their default slot priority of 1, and the less critical server modules are changed to lower priority value of 2 or higher, the priority 1 server modules is powered on first. These higher priority servers get their maximum power allocation, while lower priority servers may be not be allocated enough power to run at their maximum performance or they may not even power on at all, depending on how low the system input power cap is set and the server power requirements.

If an administrator manually powers on the low priority server modules before the higher priority ones, then the low priority server modules are the first modules to have their power allocation lowered down to the minimum value, in order to accommodate the higher priority servers. Therefore, after the available power for allocation is exhausted, CMC reclaims power from lower or equal priority servers until they are at their minimum power level.

NOTE: I/O module, fans, and mainboard, physical disk drives, storage adapters are given the highest priority. CMC reclaims power only from the lower-priority devices to meet the power needs of a higher-priority device or server.

Assigning Priority Levels To Servers

When additional power is required, server priority levels determine which servers the CMC draws power from.

- NOTE: The priority you assign to a server is linked to a server's slot and not to the server itself. If you move the server to a new slot, you must re-configure the priority for the new slot location.
- i NOTE: To perform power management actions, you must have the Chassis Configuration Administrator privilege.

Assigning Priority Levels To Servers Using CMC Web Interface

To assign priority levels:

- 1. In the left pane, click **Server Overview** > **Power** > **Priority**. The **Server Priority** page lists all the servers in the chassis.
- 2. From the **Priority** drop-down menu, select a priority level (1–9, where 1 is the highest priority) for one, multiple, or all servers. The default value is 1. You can assign the same priority level to multiple servers.
- 3. Click **Apply** to save your changes.

Assigning Priority Levels To Servers Using RACADM

Open a serial/Telnet/SSH text console to CMC, log in, and type:

 $\verb|racadm| config -g cfgServerInfo -o cfgServerPriority -i <|slot number> <|priority level>|$

where <slot number> (1-4) refers to the location of the server, and <pri>priority level> is a value between 1-9.

For example, to set the priority level to 1 for the server in slot 4, type the following command:

racadm config -g cfgServerInfo -o cfgServerPriority -i 4 1

Viewing Power Consumption Status

CMC provides the actual input power consumption for the entire system.

Viewing Power Consumption Status Using CMC Web Interface

In the left pane, click **Chassis Overview** > **Power** > **Power Monitoring**. The Power Monitoring page displays the power health, system power status, real-time power statistics, and real-time energy statistics. For more information, see the *Online Help*.

i NOTE: You can also view the power redundancy status under Power Supplies.

Viewing Power Consumption Status Using RACADM

To view power consumption status using RACADM:

Open a serial/Telnet/SSH text console to CMC, log in, and type:

racadm getpminfo

AC Power Recovery

If the AC power supply of a system is disrupted, the chassis is restored to the previous power state before the AC power loss. The restoration to the previous power state is the default behavior. The following factors could cause the disruption:

- power outage
- power cables pulled from the power supply units (PSUs)
- power distribution unit (PDU) outage

If the **Budget/Redundancy Configuration > Disable AC Power Recovery** option is selected, the chassis remains powered off after the AC recovery.

In case, the blade servers are not configured to automatic power-up, you may have to power them on manually.

Viewing Power Budget Status Using CMC Web Interface

To view power budget status using CMC Web interface, in the left pane go to **Chassis Overview** and click **Power > Budget Status**. The **Power Budget Status** page displays the system power policy configuration, power budget details, budget allocated for server modules, and chassis power supply details. For more information, see the *Online Help*.

Viewing Power Budget Status Using RACADM

Open a serial/Telnet/SSH text console to CMC, log in, and type:

racadm getpbinfo

For more information about **getpbinfo**, including output details, see the **getpbinfo** command section in the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide*.

Redundancy Status and Overall Power Health

The redundancy status is a factor in determining the overall power health. When the power redundancy policy is set, for example, to Grid Redundancy and the redundancy status indicates that the system is operating with redundancy, the overall power health is typically **OK**. If the PSU installed on a chassis fails owing to some reason, the overall power health status of the chassis is displayed as **Non-Critical**. However, if the conditions for operating with Grid redundancy cannot be met, the redundancy status is **No**, and the overall power health is **Critical**. This is because the system is not able to operate in accordance with the configured redundancy policy.

NOTE: CMC does not perform a pre-check of these conditions when you change the redundancy policy to or from Grid redundancy. So, configuring the redundancy policy may immediately result in redundancy lost or a regained condition.

Power Management After PSU Failure

When an insufficient-power event occurs, such as a PSU failure, CMC reduces power supply to the servers. After reducing the power, CMC reevaluates the power needs of the chassis. If power requirements are still not met, CMC turns off the lower priority servers. However, this is done on the basis of power redundancy policy that you set on your CMC. A redundant server can tolerate the loss of power without impacting the performance of the servers.

Power for higher priority servers is restored incrementally, while power needs remain within the power budget. To set the redundancy policy, see Configuring Power Budget and Redundancy.

Power Management After Removing PSU

CMC may begin conserving power when you remove a PSU or a PSU AC cord. CMC decreases power to the lower priority servers until power allocation is supported by the remaining PSUs in the chassis. If you remove more than one PSU, CMC again evaluates the power requirements when the second PSU is removed to determine the firmware response. If power requirements are still not met, CMC may turn off the low-priority servers.

Limits

- CMC does not support *automated* power-down of a low-priority server to allow turning-on of a higher priority server; however, you can perform user-initiated turn-offs.
- Changes to the PSU redundancy policy are limited by the number of PSUs in the chassis. You can select any of the two PSU redundancy configuration settings listed in Default Redundancy Configuration.

New Server Engagement Policy

If a new server that is turned on exceeds the power available for the chassis, CMC may decrease the power to the low-priority servers. This could happen if the administrator has configured a power limit for the chassis that is below what would be required for full power allocation to the servers, or if insufficient power is available in case higher power requirements by all servers in the chassis. If enough power cannot be freed by reducing the allocated power of the low-priority servers, the new server is not allowed to turn on.

This occurs if the administrator had configured power limit for the chassis lower than the full power allocation to the servers or if insufficient power is available to servers requiring high power.

The following table provides the actions taken by CMC, when a new server is powered on in the scenario described earlier.

Table 37. CMC Response When a Server Power-On is Attempted

Worst Case Power is Available	CMC Response	Server Power On
Yes	No power conservation is required	Allowed
No	Perform power conservation: • Power required for new server is available • Power required for new server is not available	Allowed Not Allowed

If a PSU stops functioning, it results in a non-critical health state and a PSU failure event is generated. The removal of a PSU results in a PSU removal event.

If either event results in a loss of redundancy, on the basis of power allocations, a loss of redundancy event is generated.

If the subsequent power capacity or the user power capacity is greater than the server allocations, servers have degraded performance or, in an extrement case, servers may be turned off. Both conditions are in reverse-priority order; that is, the low-priority servers are turned off first.

The following table provides the firmware response to a PSU turnoff, or removal as it applies to various PSU redundancy configurations.

Table 38. Chassis Impact from PSU Failure or Removal

PSU Configuration	Dynamic PSU Engagement	Firmware Response
Grid Redundancy	Disabled	CMC alerts you about loss of Grid redundancy.
Power Supply Redundancy	Disabled	CMC alerts you about loss of power supply redundancy.
Grid Redundancy	Enabled	CMC alerts you about loss of Grid Redundancy. PSUs in standby mode (if any) are turned on to compensate for power budget lost from the PSU failure or removal.
Power Supply Redundancy	Enabled	CMC alerts you of loss of Power Supply Redundancy. PSUs in standby mode (if any) are turned on to compensate for power budget lost from PSU failure or removal.

Power Supply and Redundancy Policy Changes in System Event Log

Changes in the power supply state and power redundancy policy are recorded as events. Events related to the power supply that record entries in the system event log (SEL) are power supply insertion and removal, power supply input insertion and removal, and power supply output assertion and de-assertion.

The following table lists the SEL entries that are related to power supply changes:

Table 39. SEL Events for Power Supply Changes

Power Supply Event	System Event Log (SEL) Entry	
Insertion	Power supply is present.	
Removal	Power supply is absent.	
AC input received	The power input for power supply has been restored.	
AC input lost	The power input for power supply is lost.	
DC output produced	Power supply is operating normally.	
DC output lost	Power supply failed.	

Events related to changes in the power redundancy status that record entries in the SEL are redundancy loss and redundancy regain for the modular enclosure that is configured for either an **Grid Redundancy** power policy or **Power Supply Redundancy** power policy. The following table lists the SEL entries that are related to power redundancy policy changes.

Table 40. SEL Events for Power Redundancy Policy Changes

Power Policy Event	System Event Log (SEL) Entry
Redundancy lost	Power supply redundancy is lost.
Redundancy regained	The power supplies are redundant.

Configuring power budget and redundancy

You can configure the power budget, redundancy, and dynamic power of the entire chassis (chassis, servers, I/O module, KVM, CMC, and power supplies), which uses four power supply units (PSUs). The power management service optimizes power consumption and reallocates power to different modules based on the requirement.

You can configure the following:

- System Input Power Cap
- Redundancy Policy
- Enable Dynamic Power Supply Engagement
- Disable Chassis Power Button
- Max Power Conservation Mode
- Remote Power Logging
- Remote Power Logging Interval
- Server Based Power Management
- Disable AC Power Recovery

Power Conservation and Power Budget

CMC conserves power when the user-configured maximum power limit is reached. When the demand for power exceeds the user configured System Input Power Cap, CMC reduces power to servers in reverse-priority order to free power for higher priority servers and other modules in the chassis.

If all or multiple slots in the chassis are configured with the same priority level, CMC decreases power to servers in the order of increasing slot number. For example, if the servers in slots 1 and 2 have the same priority level, the power for the server in slot 1 is decreased before that of the server in slot 2.

NOTE: You can assign a priority level to each server in the chassis assigning a number from 1 through 9 for each server. The default priority level for all servers is 1. The lower the number, the higher the priority level.

The power budget is limited to a maximum of whichever set of two PSUs that is the weakest. If you attempt to set an AC power budget value that exceeds the *System Input Power Cap* value, CMC displays a message. The power budget is limited to 4800W.

Maximum Power Conservation Mode

This is enabled for Grid Redundancy or PSU Redundancy modes. CMC performs maximum power conservation when:

- Maximum conservation mode is enabled.
- An automated command line script, issued by a UPS device, enables maximum conservation mode.

In maximum power conservation mode, all servers start functioning at their minimum power levels, and all subsequent server power allocation requests are denied. In this mode, the performance of powered on servers may be degraded. Additional servers cannot be powered on, regardless of server priority.

The system is restored to full performance when the maximum conservation mode is cleared.

NOTE: If the Maximum Power Conversation Mode (MPCM) is enabled on the chassis, all power requests from a blade server are denied. The blade server is not powered on if there is any action in the iDRAC or blade server requiring the host to start the power cycle.

Server Power Reduction to Maintain Power Budget

CMC reduces power allocations of low-priority servers, when additional power is required to maintain the system power consumption within the user-configured *System Input Power Cap*. For example, when a new server is engaged, CMC may decrease power to low-priority servers to allow more power for the new server. If the amount of power is still insufficient after reducing power allocations of the lower priority servers, CMC lowers the performance of servers until sufficient power is freed to power the new server.

CMC reduces server power allocation in two cases:

- Overall power consumption exceeds the configurable System Input Power Cap.
- A power failure occurs in a non-redundant configuration.

110V PSUs AC Operation

By default, the 110V PSU AC Operation feature is available. However, a combination of 110V and 220V operation is not supported. If CMC detects that both voltages are input, one voltage value is selected and those power supplies connected to the other voltage level are turned off and indicated as not functioning.

Remote Logging

Power consumption can be reported to a remote syslog server. Total chassis power consumption, minimum, maximum, and average power consumption over a collection period can be logged. For more information about enabling this feature and configuring the collection or logging interval, see Managing and Monitoring Power.

External Power Management

CMC power management is optionally controlled by the OpenManage Power Center (OMPC). For more information, see the OMPC User's Guide.

When external power management is enabled, OMPC manages:

- Server power of supported VRTX servers
- Server priority of supported VRTX servers
- System input power capacity
- Maximum power conservation mode

The CMC continues to maintain or manage:

- Redundancy policy
- Remote power logging
- Server performance over power redundancy
- Dynamic power supply engagement

OPMC then manages prioritization and power of supported VRTX server nodes in the chassis from the budget available after allocation of power to chassis infrastructure and prior generation server nodes. Remote power logging is unaffected by external power management.

After the Server Based Power Management Mode is enabled, the chassis is prepared for PM3 management. All supported VRTX server priorities are set to 1 (High). PM3 manages the server power and priorities directly. Since PM3 controls compatible server power allocations, CMC no longer controls the Maximum Power Conservation Mode. Hence, this selection is disabled.

When the **Maximum Power Conservation Mode** is enabled, the CMC sets the System Input Power Capacity to the maximum that the chassis can handle. CMC does not allow power to exceed the highest capacity. However, PM3 handles all other power capacity limitations.

When PM3 management of power is disabled, the CMC reverts to the server priority settings before the external management was enabled.

NOTE: When PM3 management is disabled, CMC does not revert to the earlier setting of the maximum chassis power. See the **CMC log** for the earlier setting to manually restore the value.

Configuring Power Budget and Redundancy Using CMC Web Interface

i NOTE: To perform power management actions, you must have the Chassis Configuration Administrator privilege.

To configure power budget:

- 1. In the left pane, click Chassis Overview > Power > Configuration.
- 2. On the **Budget/Redundancy Configuration** page, select any or all of the following properties as appropriate. For information about the field descriptions, see the *Online Help*.
 - Enable Server-Based Power Management
 - System Input Power Cap
 - Redundancy Policy
 - Enable Dynamic Power Supply Engagement
 - Disable Chassis Power Button
 - Max Power Conservation Mode
 - Enable Remote Power Logging
 - Remote Power Logging Interval
- 3. Click Apply to save the changes.

Configuring Power Budget and Redundancy Using RACADM

i NOTE: To perform power management actions, you must have the Chassis Configuration Administrator privilege.

To enable and set the redundancy policy:

- 1. Open a serial/Telnet/SSH text console to CMC and log in.
- 2. Set properties as needed:
 - To select a redundancy policy, type:

```
racadm config -g cfgChassisPower -o
cfgChassisRedundancyPolicy <value>
```

where <value> is 1 (Grid Redundancy), and 2 (Power Supply Redundancy). The default value is 2.

For example, the following command sets the redundancy policy to 1:

```
racadm config -g cfgChassisPower -o cfgChassisRedundancyPolicy 1
```

To set the power budget value, type:

```
racadm config -g cfgChassisPower -o cfgChassisPowerCap < value>
```

where < value> is a number between 938 W - 4800 W, representing the maximum power limit in Watt. The default is 4800.

For example, the following command sets the maximum power budget to 4800 Watt:

```
racadm config -g cfgChassisPower -o cfgChassisPowerCap 4800
```

• To enable or disable dynamic PSU engagement, type:

```
racadm config -g cfgChassisPower -o
cfgChassisDynamicPSUEngagementEnable <value>
```

where $<\!\!\mathit{value}\!\!>$ is 0 (disable), 1 (enable). The default is 0.

For example, the following command disables dynamic PSU engagement:

```
racadm config -g cfgChassisPower -o
cfgChassisDynamicPSUEngagementEnable 0
```

• To enable the maximum power consumption mode, type:

```
racadm config -g cfgChassisPower -o cfgChassisMaxPowerConservationMode 1
```

• To restore normal operation, type:

```
racadm config -g cfgChassisPower -o cfgChassisMaxPowerConservationMode 0
```

• To enable the power remote logging feature, enter the following command:

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

To specify the desired logging interval, enter the following command:

```
\begin{tabular}{ll} racadm & config -g & cfgRemoteHosts -o \\ cfgRhostsSyslogPowerLoggingInterval & n \end{tabular}
```

where n is 1-1440 minutes.

To determine if the power remote logging feature is enabled, enter the following command:

```
racadm getconfig -g cfgRemoteHosts -o
cfgRhostsSyslogPowerLoggingEnabled
```

• To determine the power remote logging interval, enter the following command:

```
racadm getconfig -g cfgRemoteHosts -o
cfgRhostsSyslogPowerLoggingInterval
```

The power remote logging feature is dependent on previously configured remote syslog hosts having been . Logging to one or more remote syslog hosts must be enabled, otherwise power consumption is logged. This can be done either through the web GUI or the RACADM CLI. For more information, see the remote syslog configuration instructions.

• To enable remote power management by Open Manage Power Center (OPMC), type:

```
racadm config -g cfgChassisPower -o
cfgChassisServerBasedPowerMgmtMode 1
```

To restore CMC power management, type:

```
racadm config -g cfgChassisPower -o
cfgChassisServerBasedPowerMgmtMode 0
```

For information about RACADM commands for chassis power, see the **config**, **getconfig**, **getpbinfo**, and **cfgChassisPower** sections in the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide*.

Executing Power Control Operations

You can execute the following power control operation for the chassis, servers, and IOM.

i NOTE: Power control operations affect the entire chassis.

Executing Power Control Operations on the Chassis

CMC enables you to remotely perform several power management actions, such as an orderly shutdown on the entire chassis (chassis, servers, IOM, and PSUs).

iNOTE: To perform power management actions, you must have the Chassis Configuration Administrator privilege.

Executing Power Control Operations on the Chassis Using Web Interface

To execute power control operations on the chassis using the CMC web interface:

- In the left pane, click Chassis Overview > Power > Control.
 The Chassis Power Control page is displayed.
- 2. Select one of the following power control operations.

For information about each option, see the Online Help.

- Power On System
- Power Off System
- Power Cycle System (cold boot)
- Reset CMC (warm boot)
- Non-Graceful Shutdown
- 3. Click Apply.

A dialog box appears asking you for a confirmation.

4. Click **OK** to perform the power management action (for example, cause the system to reset).

Executing Power Control Operations on the Chassis Using RACADM

Open a serial/Telnet/SSH text console to CMC, log in, and type:

racadm chassisaction -m chassis <action>

where <action> is powerup, powerdown, powercycle, nongraceshutdown, or reset.

Executing Power Control Operations on a Server

You can remotely perform power management actions for multiple servers at a time or an individual server in the chassis.

i NOTE: To perform power management actions, you must have the Chassis Configuration Administrator privilege.

Executing Power Control Operations for Multiple Servers Using CMC Web Interface

To execute power control operation for multiple servers using the Web interface:

- 1. In the left pane, click **Server Overview** > **Power**.
 - The Power Control page is displayed.
- 2. In the **Operations** column, from the drop-down menu, select one of the following power control operations for the required servers:
 - No Operation
 - Power On Server
 - Power Off Server
 - Graceful Shutdown
 - Reset Server (warm boot)Power Cycle Server (cold boot)
 - For information about the options, see the Online Help.
- 3. Click Apply.
 - A dialog box appears requesting for confirmation.
- 4. Click **OK** to perform the power management action (for example, reset the server).

Executing Power Control Operations on the IOM

You can remotely reset or turn on an IOM.

i NOTE: To perform power management actions, you must have the Chassis Configuration Administrator privilege.

Executing Power Control Operations on IOM Using CMC Web Interface

To execute power control operations on the I/O Module:

- 1. In the left pane, click Chassis Overview > I/O Module Overview > Power.
- 2. On the **Power Control** page, for the IOM, from the drop-down menu, select the operation you want to execute (power cycle).
- 3. Click Apply.

Executing Power Control Operations on the IOM Using RACADM

To execute power control operations on the IOM using RACADM, open a serial/Telnet/SSH text console to CMC, log in, and type:

racadm chassisaction -m switch <action>

where <action> indicates the operation you want to execute: power cycle.

Managing Chassis Storage

On the Dell PowerEdge VRTX, you can perform the following operations:

- View the status of physical disks drives and storage controllers.
- View the properties of controllers, physical disk drives, virtual disks, and enclosures.
- Set up controllers, physical disk drives, and virtual disks.
- Assign virtual adapters.
- Troubleshoot controller, physical disk drives, and virtual disks.
- Update storage components.
- Use the shared storage controllers in fault-tolerant mode
- Enabling or disabling Shared PERC8 (Integrated 2)
- NOTE: Initialize fast or initialize full is not displayed when virtual disks are initially created.

Topics:

- Viewing Status of the Storage Components
- Viewing the Storage Topology
- Viewing Fault-tolerant Troubleshooting Information of SPERC Using CMC Web Interface
- Assigning Virtual Adapters To Slots Using CMC Web Interface
- Fault-Tolerance in Storage Controllers
- Security Key Mismatch
- Viewing Controller Properties Using CMC Web Interface
- Viewing Controller Properties Using RACADM
- Configuring Storage Controller Settings
- Shared PERC Controllers
- Enabling or Disabling RAID Controller Using CMC Web Interface
- Enabling or Disabling RAID Controller Using RACADM
- Enabling or disabling fault tolerance of external RAID controller using RACADM
- Viewing Physical Disk Properties Using the CMC Web Interface
- Viewing Physical Disk Drives Properties Using RACADM
- Identifying Physical Disks and Virtual Disks
- Assigning Global Hot Spares Using CMC Web Interface
- Assigning Global Hot Spares Using RACADM
- Recovering Physical Disks
- Viewing Virtual Disk Properties Using CMC Web Interface
- Viewing Virtual Disk Properties Using RACADM
- Creating Virtual Disk Using CMC Web Interface
- Managing Encryption Keys
- Encrypting Virtual Disks
- Unlocking Foreign Configuration
- Cryptographic Erase
- Applying Virtual Adapter Access Policy To Virtual Disks
- Modifying Virtual Disk Properties Using CMC Web Interface
- Enclosure Management Module
- Viewing Enclosure Properties Using CMC Web Interface

Viewing Status of the Storage Components

To view the status of the storage components:

1. In the left pane, click Chassis Overview > Storage > Properties > Storage Overview.

- 2. On the Storage Overview page, you can:
 - View the graphic summary of the physical disk drives installed in the chassis and their status.
 - View the summary of all the storage components with links to their respective pages.
 - View the used capacity and total capacity of the storage.
 - View controller information.
 - NOTE: In case of a fault-tolerant controller, the name format is: Shared < PERC number > (Integrated < number >). For example, the active controller is Shared PERC8 (Integrated 1) and the peer controller is Shared PERC8 (Integrated 2).
 - NOTE: If the secondary PERC is disabled, the name is displayed as Disabled PERC (Integrated 2).
 - View recently-logged storage events.
 - i NOTE: For more information, see the Online Help.

Viewing the Storage Topology

To view the Storage topology:

- 1. In the left pane, click Chassis Overview > Storage > Properties > Topology.
- 2. On the **Topology** page, click the **<controller name>** to view the respective pages.
 - NOTE: You can view the name of the controller that is active in controlling the storage devices associated with this CMC and also the passive controller acting as a stand-by.
- 3. Under each installed controllers, click the links **View Virtual Disks**, **<enclosure name>**, and **View Physical Disks** to open the respective pages.

Viewing Fault-tolerant Troubleshooting Information of SPERC Using CMC Web Interface

To view the attributes that indicate the correct functioning of fault-tolerant features of a SPERC:

- In the left pane, click Chassis Overview > Storage > Troubleshooting > Setup Troubleshooting.
 The Storage Setup Troubleshooting page is displayed.
- 2. On the Storage Setup Troubleshooting page, you can:
 - Click to view the following attributes when the integrated controller is in fault-tolerant mode:
 - o Two Shared PERCs detected.
 - o Two expanders detected
 - o Shared PERCs and expanders correctly cabled
 - Correct Firmware on Shared PERCs
 - o Correct Firmware on Expanders
 - o Correct Firmware on Chassis Infrastructure
 - o Shared PERCs have the same settings: Indicates whether or not SPERCs have the same settings.
 - ullet Click ullet to view the following attributes when the integrated controller is not in fault-tolerant mode:
 - o One Shared PERC detected
 - One expanders detected
 - o Shared PERC and expanders correctly cabled
 - Click to view the following attributes when the external controller is in fault tolerant:
 - Two Shared PERCs detected
 - Shared PERCs are installed in different fabrics
 - o Shared PERCs and EMMs are connected correctly
 - Correct firmware on Shared PERCs
 - Shared PERCs have the same settings
 - Click to view the following attributes when the external controller is not in fault-tolerant mode:

- o One Shared PERC detected
- One expanders detected
- Shared PERC and expanders correctly cabled
- View the status of each attribute that indicates if the fault-tolerant criteria is fulfilled.
 - NOTE: If the attribute in a fault-tolerant environment is not matching the criterion, then an **Update Now** option is displayed for that attribute.
- NOTE: A Learn How option is displayed against some of the attributes. For more information about the attribute, click Learn How.
- 3. To fulfill a criterion for an attribute, click **Update Now**.
 - The **Storage Component Update** page is displayed, which allows you to update the required storage component to fulfill the criterion for the attribute.

Assigning Virtual Adapters To Slots Using CMC Web Interface

Using the Virtual Adapter feature, you can share the installed storage with the four servers. You can map a virtual disk to a server slot by first mapping a virtual disk to a virtual adapter (VA), and then mapping a Virtual Adapter (VA) to a server slot.

- Before assigning a VA to a server slot, ensure that:
 - The server slot is empty, or the server in the slot is turned off.
 - The VA is unmapped from a server or a slot.
 - o All affected servers are powered off.
- Virtual disks are created and they are assigned as Virtual Adapter 1, Virtual Adapter 2, Virtual Adapter 3, or Virtual Adapter 4. For more information, see Applying Virtual Adapter Access Policy to Virtual Disks.
- (i) NOTE:
 - You can map only one virtual adapter to one server at a time.
 - Without an appropriate license, you can unmap a VA-server assignment, or map the VA to the default serve, only.
 - The default mapping is VA1-Server Slot 1, VA2-Server Slot 2, VA3-Server Slot 3, and VA4-Server Slot 4.
 - If full height server is inserted, then the upper slot has the VA mapped to it while the bottom slot is still unmapped. For example, a full height in slot 1 has VA1 assigned to slot 1 and VA3 is still unmapped.
 - If the system has an enterprise license, then you can assign any one of the four VAs to a server slot. However, you can still map one virtual adapter to one server at a time.
 - Virtual adapter rules are applied to the external and integrated shared storage adapters.

To map or unmap a virtual adapter from a server slot:

- In the left pane, click Chassis Overview > Storage > Setup > Virtualization.
 The Storage Virtualization page is displayed.
- 2. To select the required type of assignment, from the Assignment Mode: Virtual Disks to Virtual Adapters table, select:
 - **Single Assignment** Select to assign one virtual disk to one virtual adapter.
 - **Multiple Assignment** Select to assign a virtual disk to multiple virtual adapters. Read the on-screen instructions before selecting this option.
 - NOTE: Select the **Multiple Assignment** mode only when the servers have Cluster Services installed on them. Use of this mode without Cluster Services may lead to corrupted or lost data.
 - NOTE: You can assign a virtual disk to multiple virtual adapters from CMC CLI even when the **Assignment Mode** is set to **Single Assignment** in the CMC web interface.
- 3. In the **Virtual Adapters Mapped** table, from the **Action** drop-down list, select one of the following options, and then click **Apply**.
 - <Slot #> Select the slot to which the VA must be assigned.
 - **Unmap** Select to remove the VA assignment to a slot.

The VA is mapped or unmapped from the selected server slot, based on the selected action.

NOTE: Consider a VA assigned to server in the lower slot (3 or 4). When a half height server (slot 3 or 4) is replaced with a full height server, the full height server does not access the VA assigned to lower slots. Inserting a half height server again, provides access to the VA.

Map or un-map a PERC Virtual Controller to blade:

- Each External Shared PERC 8 card has four virtual adapters (VA). If one or two External Shared PERC 8 cards are
 present in the system, then in shared mode you can map or un-map one of the four virtual adapters'.
- If an external PCIE slot is occupied by shared adapter, the virtual adapter mapping can get the current details or information for VA mapping of the shared storage VA pool.
- Shared device is not supported when the external PCIE slot is occupied by a shared adapter. Using the shared adapter, you can support shared device by changing the shared storage VA pool.

Fault-Tolerance in Storage Controllers

High Availability (HA) in storage enables availability of multiple integrated components and multiple access points to storage resources. In case a storage component stops functioning, the server is supported by a second critical component or path to the available data. High Availability only minimizes downtime by restoring services behind the scenes, in most cases before the non-functionality is visible, but does not eliminate downtime. Fault Tolerance (FT) makes use of redundant components within a storage system, which are configured to behave as backup components and are kept in standby mode. Storage Controllers in fault-tolerant mode prevent disruption of storage services and automatically take over the services of a component that has stopped functioning. Performance remains consistent throughout this failover process since the redundant components(controllers) are not used during normal operating conditions.

High Availability with fault-tolerance provides the following benefits:

- Provides uptime for all storage applications even when a controller stops functioning.
- Provides access to critical functions of the chassis at all times.
- Enables server to handle situations when controller stops functioning are becomes faulty.
- Makes use of component redundancy

Using the fault-tolerant feature of controllers, you can manage the tasks associated with shared storage that are achieved by having an active and passive (peer) controller. Active controller is the controller that is active and monitors all the storage-related processes. Working status of both controllers are communicated between controllers so that when an active controller stops functioning, the passive controller acting as a peer hot-spare, takes over in a seamless manner.

- NOTE: CMC displays fault-tolerant data for Shared PERC 8's with SR-IOV enabled firmware. If a non-SR-IOV card is attached to the shared storage slots the card does not power on and an alert is generated.
- NOTE: Operations such as resetting CMC, which reset the CMC configuration, reset the external fault-tolerant configuration. As a result, the PERC mode changes to "Safe mode". Disable the fault-tolerance in the external PERC.

Security Key Mismatch

You can create a security key on a controller using an **Encryption keyID** and a **Passphrase**. The controller compares only the **Passphrase** used while creating security key to identify whether the two controllers have the same security keys. Therefore, two controllers joining a cluster are fault-tolerant even if they have different **Encryption Key ID** as long as they have the same passphrase.

If a security key mismatch is detected between two peer controllers, the fault-tolerant mode changes to 'Degraded'. A critical alert is displayed on the **Chassis Health** page and monitoring may not display proper drive association.

If a security key mismatch is detected, resolve the key mismatch by creating, modifying, or deleting the security key on one of the controllers, before performing any other storage security operation on the controller. Power cycle the chassis after resolving the mismatch. Before combining two non-high availability controllers, modify the keys so that they match. This action facilitates import of secure drives associated with each controller joining the cluster.

For external controllers, modify the keys so that they match before cabling them for fault tolerance. The modification of security keys facilitates import of secure drives associated with each controller joining the cluster.

Resolving Security Key Mismatch Using CMC Web Interface

To resolve the security key mismatch using the CMC web interface:

- 1. Turn off the server modules.
- 2. Click Server Overview > Power > Control > Power Off Server.
- 3. Modify the security key on one or both of the existing non-fault-tolerant controllers so that the keys match.
- 4. Power cycle the chassis.
- 5. Verify if the controllers have matching keys.

Viewing Controller Properties Using CMC Web Interface

To view the controller properties:

- 1. In the left pane, click Chassis Overview > Controller.
- 2. On the **Controllers** page, under the **Controllers** section, you can see the basic properties of the controller. However, to view the advanced properties, click the +.
 - NOTE: If the Controllers are in fault-tolerant mode, then the following information regarding the fault-tolerant status and mode is also displayed:
 - Fault Tolerant Mode Shared, Active/Passive
 - Fault Tolerant Status Healthy/Normal, or Lost/Degraded
 - Peer Controller Indicates the name of controller that acts as the peer (stand-by) in case of a fault-tolerant mode supported by two controllers
 - NOTE: If the peer controller is disabled, the name is displayed as **Disabled PERC (Integrated 2)** or **Disabled PERC(SPERC Slot 6)** and the Status is displayed as **Unknown**, which implies that the peer controller is turned off.

For more information about Controllers, see the Online Help.

Viewing Controller Properties Using RACADM

To view controller properties using RACADM, run the command racadm raid get controllers -o

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Importing or Clearing Foreign Configuration

A foreign disk must be inserted into the chassis.

To import or clear the foreign configuration:

- 1. In the left pane, click Chassis Overview > Storage > Controllers > Setup.
- 2. On the Controller Setup page, in the Foreign Configuration section, for the respective controller, click:
 - Clear Foreign Configuration to clear the existing configuration of the disk.
 - Import/Recover to import the disk with the foreign configuration.

Configuring Storage Controller Settings

You can modify existing properties of a storage controller or configure the properties of a newly-installed storage controller.

Configuring Storage Controller Settings Using CMC Web Interface

Make sure at least one storage controller is installed in the chassis.

Toconfigure the storage controller settings:

- 1. In the CMC Web interface, go to Chassis Overview > Storage > Controllers > Setup.
- 2. On the Controller Setup page, from the Controller drop-down menu, select the controller.
 - NOTE: Note the following:
 - If the storage controllers are in fault-tolerant mode, and if both have the same firmware version, then both the controllers are displayed as a single device in the drop-down menu. For example, Shared PERC8 (Integrated 1) or Shared PERC8 (Integrated 2) or Shared PERC8 (SPERC Slot 5) or Shared PERC8 (SPERC Slot 6). If the settings of two controllers are different, the **Settings Incompatible** message is displayed. You can set the properties of fault-tolerant controllers so that the properties are same on both the controllers. Controllers in this mode cannot have separate properties.
 - If a second storage controller with a different firmware version is installed, then the controllers are displayed as two different components in the drop-down menu. For example, Shared PERC8 (Integrated 1), Shared PERC8 (Integrated 2), Shared PERC8 (SPERC Slot 5) and Shared PERC8 (SPERC Slot 6).

The attribute values for the selected controller are updated in the table.

- 3. Type or select appropriate data, and then click Apply.
 - NOTE: For information about the attributes and other field descriptions, see Online Help.

The newly set properties are applied to the selected controllers and the **Current Value** field displays the updated values for the attributes.

Configuring Storage Controller Settings Using RACADM

To set up the storage controller by running a RACADM command, use the following syntax.

```
racadm raid ctrlprop:RAID.ChassisIntegrated.1-1 [-rebuild <value>] [-bgi <value>] [-
reconstruct <value>] [-checkconsistency <value>] [-ccmode {abortonerror | normal}] [-
copybackmode {off | on | onwithsmart}] [-lb {auto | disabled}] [-prunconfigured {yes |
no}]
```

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Shared PERC Controllers

For systems with two integrated Shared PERC installed, you can change the operation mode from **Fault Tolerant** mode to **Non Fault Tolerant** mode, or contrariwise using the Web Interface or the RACADM CLI by enabling or disabling the second internal shared PERC 8 controller.

For internal shared PERC8 controller, you can disable the second integrated controller. After the second integrated controller is disabled the First integrated controller is not in fault-tolerant mode. When the second integrated controller is enabled, the two integrated controllers are in fault-tolerant mode, by default. The second integrated controller can be disabled by using racadm raid disableperc:RAID.ChassisIntegrated.2-1 command.

For external enclosures, both the External Shared PERC 8 card in slot 5 and slot 6 can be disabled using racadm raid disableperc: RAID.ChassisSlot.5-1 and racadm raid disableperc: RAID.ChassisSlot.6-1 command respectively.

From the RACADM Command Line Interface, run the racadm raid get controllers command to list the number of Shared PERC controllers on your system. If the command lists only RAID.ChassisIntegrated.1-1, your system has a single Shared PERC controller. If the command lists RAID.ChassisIntegrated.1-1, RAID.ChassisIntegrated.2-1 your system has two Shared PERC controllers.

Second Integrated Shared PERC 8 and External Shared PERC 8 cards in slot 5 and slot 6 can be enabled or disabled.

To change the operation mode using the CMC Web Interface, go to the **Controllers Troubleshooting** page by navigating to **Chassis Overview** > **Storage Controllers**, in the left pane and select the **Disable Raid Controller** or **Enable Raid Controller** option.

To change the operational mode using the RACADM CLI:

- Run the racadm raid enableperc:RAID.ChassisIntegrated.2-1 command to enable the Integrated 2 shared PERC 8 and Fault Tolerant mode, if second integrated shared PERC8 is disabled.
- Run the racadm raid enableperc: RAID. ChassisSlot. 6-1 command to enable external Shared PERC8 in slot 6.
- Run the racadm raid disableperc:RAID.ChassisIntegrated.2-1 command to disable **Second Integrated** shared PERC8 and Fault Tolerant mode.

(i) NOTE:

- The chassis must be turned on and all server modules must be turned off before you run the enable or disable commands. The chassis is automatically power cycled as part of this operation. After changing the Shared PERC operation mode, it is recommended to reset the CMC using the **Troubleshooting** page or the racadm racreset
- By default if second integrated PERC 8 cards are detected, then the mode displays high availability mode.
- Enabling the SPERC in external slots do not enable fault tolerance.
- To enable Fault Tolerant mode for external shared PERC8, refer *Enable or Disable Fault Tolerance of External RAID controller using RACADM* section.

Enabling or Disabling RAID Controller Using CMC Web Interface

For a VRTX chassis with two Shared PERC8 controllers, the Integrated 2 PERC adapter can be disabled or enabled when the Integrated 1 PERC adapter is active and all server modules are turned off. Both adapters must be enabled for fault-tolerance. The **Controllers Troubleshooting** page allows you to enable or disable the peer controller.

- (i) NOTE: To prevent loss of data, before performing controller enable or disable operations:
 - Complete all data operations such as Rebuild or Copy Back.
 - Make sure that the data volumes are in optimal state.
- NOTE: While enabling the second PERC adapter, a warning message is displayed and the fault-tolerant status is degraded if:
 - Any of the PERC adapter settings are changed.
 - Firmware is updated.

Make sure that the firmware and the settings of the shared PERC match to set the fault tolerant system configuration to Fault Tolerant mode.

You can disable a peer controller only if:

- All the servers in the chassis are turned off.
- The Integrated 1 PERC is currently the active controller.
 - NOTE: If the Integrated 1 PERC is not currently the active controller, then perform a chassis power cycle to make this the active controller.
- Both CMCs have the same firmware version that supports this feature.
- NOTE: After disabling Integrated 2 PERC, to replace a CMC card, it is recommended to update the CMC card with firmware version 1.35 or later, before the card is assigned to be the active CMC controller in the system. A message is displayed before you perform this action.

To enable or disable a peer controller in fault-tolerant mode using the CMC Web interface:

- 1. In the left pane, click Chassis Overview > Storage > Controllers > Troubleshooting.
- 2. On the **Controller Troubleshooting** page, from the **Actions** drop-down menu for Integrated 2 PERC, select one of the following, and then click **Apply**.

- Disable RAID Controller Disables the peer controller in fault-tolerant mode.
- **Enable RAID Controller** Enables the peer controller in fault-tolerant mode. If the Integrated 2 PERC is already disabled, then the **Enable Raid Controller** option is available in the drop-down menu.
- To enable or disable External Shared PERC 8 card controllers :
 - On the **Controller Troubleshooting** page, from the **Actions** drop-down menu for External Shared PERC 8 card in slot 5 or slot 6, select one of the following, and then click **Apply**.
 - **Disable RAID Controller** Disables the RAID controller.
 - Enable RAID Controller Enables the RAID controller. If the PERC is already disabled, then Enable Raid Controller option is available in the drop-down menu.
- **Reset Configuration** Select this option to delete virtual drive and unassign all the hot spares attached to the controller. However, this only removes the disks from the configuration and does not delete any data. NOTE: Reset Configuration does not remove any foreign configurations. Use Clear Foreign Configuration to reset.
- Export TTY Log Select this option to export the TTY log on the local system. NOTE: The TTY Log collected from the controller does not contain any data from the drives. However, it may contain data such as SAS addresses.
- Enable Fault Tolerance Select this option to enable the fault-tolerance mode of the external SPERC. This action also resets the External Shared PERC 8 card.
- **Disable Fault Tolerance** Select this option to disable the fault-tolerance mode of the external SPERC. This action also resets the External Shared PERC 8 card.
 - (i) NOTE:
 - For a disabled PERC, none of the other options Reset Configuration, Export TTY Log, Discard Pinned Cache, and Disable RAID Controller are available in the drop-down menu.
 - o By default the two integrated shared storage adapters are detected with high availability mode.
 - o You must enable Fault Tolerance mode on the external shared controller after it is cabled.
 - **Enable Fault Tolerance** and **Disable Fault Tolerance** are displayed only for the External Shared PERC 8 cards. The default mode of the External Shared PERC 8 cards is the non-fault-tolerant mode.

NOTE: Enabling or disabling a peer controller initiates a chassis power cycle. The changes are reflected only after the power cycle is complete.

Enabling or Disabling RAID Controller Using RACADM

To enable a peer controller using RACADM, open a serial/Telnet/SSH text console to CMC, log in, and enter:

racadm raid enableperc:<AdapterFQDD>

To disable a peer controller, enter:

racadm raid disableperc:<AdapterFQDD>

NOTE: For more information on this feature using the RACADM interface see the RACADM Command Line Reference Guide for iDRAC and CMC.

Enabling or disabling fault tolerance of external RAID controller using RACADM

To enable the fault tolerance:

racadm raid controllers: <FQDD of External Shared PERC8> -p HighAvailabilityMode ha

To disable the fault tolerance:

racadm raid set controllers: <FQDD of External Shared PERC8> -p HighAvailabilityMode None

Viewing Physical Disk Properties Using the CMC Web Interface

Make sure that physical disks are installed on the chassis.

To view the properties of physical disk drives:

- In the left pane, go to Chassis Overview > Storage > Physical Disks. The Properties page is displayed.
- 2. To view properties of all the physical disk drives, under the **Physical Disks** section, click the ±.
 - NOTE: The following attributes are displayed for fault-tolerant mode of integrated shared adapters:
 - Active controller Shared PERC8 (Integrated 1)
 - Redundant/Failover controller Shared PERC8 (Integrated 2)

The following attributes are displayed for fault-tolerant mode of external shared adapters:

- Active controller Shared PERC8 (SPERC Slot 5)
- Redundant/Failover controller Shared PERC8 (SPERC Slot 6)

You can also use the following filters to view specific physical disk drive's properties:

- Under the Basic Physical Disks Filter option, from the Group By drop-down menu, select Virtual Disk, Controller, or Enclosure, and then click Apply.
- Click Advanced Filter, select the values for various attributes, and then click Apply.

Viewing Physical Disk Drives Properties Using RACADM

To view the properties of physical disk drives using RACADM, run the command $racadm\ raid\ get\ pdisks\ -o$

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Identifying Physical Disks and Virtual Disks

For more information about enabling or disabling the LED-blink feature, see:

- Configuring LED Blinking Using CMC Web Interface
- Configuring LED Blinking Using RACADM

Assigning Global Hot Spares Using CMC Web Interface

To assign or unassign a global hot spare:

- In the left pane, click Chassis Overview > Storage > Physical Disk > Setup.
 The Configure Physical Disks page is displayed.
- 2. Under the Configure Physical Disks section, from the Physical Disk Actions drop-down menu, select Unassigned or Global Hotspare for each of the physical disk drives, and then click Apply.
- (i) NOTE: Global hot spare assignment is allowed only if at least one virtual disk is present on the corresponding controller.

Assigning Global Hot Spares Using RACADM

To assign global hot spare using RACADM, run the command racadm \mbox{raid} hotspare: -assign \mbox{yes} -type \mbox{ghs} .

For more information about using RACADM commands, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Recovering Physical Disks

To recover a physical disks:

- 1. In the CMC Web Interface, go to Chassis Overview > Storage > Physical Disks > Setup.
- 2. On the **Setup** page, under **Recover Physical Disks** section, select the physical disk that must be recovered and from the drop-down menu, appropriately select **Rebuild Drive**, **Cancel Rebuild**, or **Force Online**, and then click **Apply**.

Viewing Virtual Disk Properties Using CMC Web Interface

Make sure that the virtual disks are created.

To view the virtual disk properties:

- 1. In the left pane, click Chassis Overview > Storage > Virtual Disks > Properties.
- 2. On the **Properties** page, under the **Virtual Disks** section, click the . You can also use the following filters to view specific virtual disk properties:
 - Under Basic Virtual Disks Filter section, from the Controller drop-down menut, select controller name, and then click Apply.
 - Click Advanced Filter, select the values for various attributes, and then click Apply.

Viewing Virtual Disk Properties Using RACADM

To view virtual disk properties using RACADM, run the command racadm raid get vdisks -o

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Creating Virtual Disk Using CMC Web Interface

By default, CMC creates virtual disks without initializing them. However, you can choose the fast-initialization option for virtual disks that are created without initialization. The fast-initialization process clears the first and last 8 MB of the virtual disk, deleting all boot records or partition information. You must have the **Chassis Configuration Administrator** privilege to perform fast-initialization.

Make sure that the physical disk is installed in the chassis.

i NOTE: Deleting a virtual disk removes the virtual disk from the controller's configuration.

To create a virtual disk:

- 1. In the left pane, click Chassis Overview > Storage > Virtual Disks > Create.
- 2. On the Create Virtual Disk page, from the RAID Level section, select the RAID level.
- 3. From the Select Physical Disks section, select the number of physical disk drives based on the RAID level selected.
- 4. In the Configure Settings section, type appropriate data, select the Initialize and Encrypt Virtual Disk options, and then click Create Virtual Disk.

CMC provides a new option, initialization, while creating virtual disks. This option allows you to create a virtual disk without fast initialization. By default, the virtual disk is created with fast initialization.

The **Initialize** option allows you to create virtual disks without initialization. This option overrides the default behavior where a fast initialization process starts when a virtual disk is created.

The Encrypt Virtual Disk option allows you to create secure virtual disks on Self-Encryption Drives (SEDs).

NOTE: The Encrypt Virtual Disk option is enabled only if the encryption key is configured for the specific controller, on the Controller Settings page.

Managing Encryption Keys

An encryption or security key, created on a controller, is used to lock or unlock access to secure virtual disks created on SEDs. You can create only one encryption key for an encryption-capable controller. You can create encryption keys by entering an encryption key identifier and passphrase, on the **Controller Setup** page. CMC also allows you to modify encryption key passphrases and delete encryption keys.

Creating Encryption Key Using CMC Web Interface

You can create encryption or security keys for controllers if the encryption key is Unconfigured.

To create an encryption key:

- 1. In the left pane, go to Storage > Controllers > Setup.
- 2. From the **Security Key** drop-down, select **Create Security Key**. A pop-up window is displayed.
- 3. Enter the security key and password and click OK.
- 4. On the Controller Setup page, click Apply.

Once the encryption key is created, the status of the Security Key changes to Enabled.

Creating Encryption Key Using RACADM

To create an encryption key by running a RACADM command, use the following syntax:

 $\label{local_radial_radial} \mbox{racadm raid createsecuritykey:RAID.ChassisIntegrated.1-1 -key < Key id > -passwd < passphrase > \\$

For more information, see the Chassis Management Controller PowerEdge VRTX RACADM Command Line Reference Guide.

Modifying Encryption Key Identifier Using CMC Web Interface

You can modify the encryption key identifier and passphrase for controllers.

To modify an encryption key identifier and passphrase:

- 1. In the left pane, go to Storage > Controllers > Setup.
- From the Security Key drop-down, select Modify Security Key. A pop-up window is displayed.
- 3. Enter the new encryption key identifier and existing and new passphrases, and click **OK**.
- 4. On the Controller Setup page, click Apply.

Modifying Encryption Identifier Key Using RACADM

To modify an encryption key identifier and passphrase by running a RACADM command, use the following syntax:

 $\label{lem:condition} $$\operatorname{racadm}$ \ raid \ modifysecuritykey: RAID. Chassis Integrated. 1-1 -key < Key id > -oldpasswd < oldpassphrase > -newpasswd < newpassphrase > $$$

For more information, see the Chassis Management Controller PowerEdge VRTX RACADM Command Line Reference Guide.

Deleting Encryption Key Using CMC Web Interface

You can only delete encryption keys for a controller only when secured virtual disks are not associated with it.

To delete an encryption key:

- 1. In the left pane, go to Storage > Controllers > Setup.
- From the Security Key drop-down, select Delete Security Key. A confirmation message is displayed.
- 3. Click OK to proceed.

After you delete the encryption key, all the SEDs that are not part of the virtual disks are secure-erased. For more information, see the *Online Help*.

Deleting Encryption Key Using RACADM

To delete an encryption key by running a RACADM command, use the following syntax:

racadm raid deletesecuritykey:RAID.ChassisIntegrated.1-1

For more information, see the Chassis Management Controller PowerEdge VRTX RACADM Command Line Reference Guide.

Encrypting Virtual Disks

You can encrypt virtual disks created on SEDs after configuring an encryption key on the controller. Whenever you perform an encryption, a message is logged in the CMC Log. You can encrypt virtual disks:

- Security key is configured on the controller.
- All the drives on the virtual disk are SEDs.

Encrypting one virtual disk enables encryption on all the virtual disks on the same disks group.

You must have the Chassis Configuration Administrator privilege to encrypt virtual disks.

Encrypting Virtual Disks Using CMC Web Interface

To encrypt an existing virtual disk:

- 1. In the left pane, click Storage > Virtual Disks > Manage.
- 2. From the Virtual Actions drop-down, select Encrypt Virtual Disk and click Apply.

(i) NOTE: The Encrypt Virtual Disk option is available only if unsecure virtual disks are configured in the SED.

Encrypting Virtual Disks Using RACADM

To encrypt virtual disks by running a RACADM command, use the following syntax:

racadm raid encryptvd:Disk.Virtual.0:RAID.ChassisIntegrated.1-1

For more information, see the Chassis Management Controller PowerEdge VRTX RACADM Command Line Reference Guide.

Unlocking Foreign Configuration

Drives which are part of secure virtual disks are called secured drives. Secured drives can be migrated from one controller to another controller. If a different encryption or security key is configured for the destination controller, the security status of these drives is displayed as 'locked' and cannot be seen as part of 'preview foreign config'. The 'Import foreign config' does not detect these foreign drives.

While running the unlock command, provide the source controller passphrase and key ID for these drives. Even after unlocking, the 'foreign controller key' still secures these drives. However, you can see these drives while searching for foreign drives in the existing 'preview foreign config'. You can import or clear the foreign configuration on these secure drives.

If foreign drives with different security keys are migrated from more than one controller, then unlock and import or clear the set of drives from one foreign controller before unlocking the drives migrated from another controller. This action ensures that unlock is not allowed on a controller, if the controller has drives that are unlocked but not imported or cleared.

Once drives are unlocked, you can import the foreign configuration using the CMC web interface or RACADM.

If the controller is power cycled after the unlock and before the import phase, the drives are locked again.

If the system has multiple foreign configurations, unlock and import each foreign configuration before unlocking the foreign configuration.

The key ID used in unlocking is used only to identify the drives with matching the key ID. After the matching drives are found, the passphrase is used for unlocking the drives.

Unlocking Foreign Configuration Using CMC Web Interface

To unlock the foreign configuration:

- 1. In the left pane, click Chassis Overview > Storage > Controllers > Setup.
- 2. Navigate to the Setup page.
- Click Click here to Unlock.The Physical Disks page is displayed.
- 4. Select the physical disks that you want to unlock.
- 5. Check if the Physical disk is associated with the key identifier.
- From the Actions drop-down, select Unlock Drive.A dialogue box is displayed prompting you to enter the security key phrase.
- 7. Enter a passphrase in the Security Key Passphrase text box.
- Re-enter the passphrase and click Unlock.
 The physical drive is unlocked and the drive does not appear in the Recover Physical Disks list.

Unlocking Foreign Configuration Using RACADM

To unlock foreign configuration by running a RACADM command, use the following syntax:

racadm raid unlock:<Controller FQDD> -key <Key id> -passwd <passphrase>

For more information, see the Chassis Management Controller PowerEdge VRTX RACADM Command Line Reference Guide.

Cryptographic Erase

You can use the cryptographic erase option to securely erase data present on secure SEDs. Secure data exists on drives even after the virtual disk is deleted and is thus exposed to threat. Cryptographic erase can be used in the following conditions:

- To erase data to retire/reuse secure drives.
- To securely erase data if secure and locked foreign configuration need not be imported.
- To recover locked drives if the passphrase is lost.

You can perform the cryptographic erase on one or more SED physical disks.

CAUTION: Performing the cryptographic erase task erases all data on the physical disk.

Performing Cryptographic Erase

If the physical disk is part of a virtual disk, remove it from the virtual disk before performing cryptographic erase.

To perform a cryptographic erase:

1. In the left pane, go to Storage > Physical Disks > Setup.

The Configure Physical Disks page is displayed.

- 2. Select the physical disk from which you want to erase the data.
- From the Physical Disk Actions drop-down, select Cryptographic Erase and click Apply.
 A message is displayed prompting you to confirm the action.
- Click Yes to proceed
 All data from the selected physical disk is removed.

Applying Virtual Adapter Access Policy To Virtual Disks

Make sure that physical disk drives are installed and the virtual disks are created.

To apply the virtual adapter access policy:

- 1. In the left pane, click Chassis Overview > Storage > Virtual Disks > Assign.
- 2. On the Assign Virtual Disks page, under the Access Policy for Virtual Adapters section, from the Virtual Adapter <number> drop-down menu, select Full Access to each physical disk drive.
- 3. Click Apply.

You can now assign virtual adapters to server slots. For more information, see the Assigning Virtual Adapters to Slots section in this User's Guide.

Modifying Virtual Disk Properties Using CMC Web Interface

To modify the virtual disk properties:

- 1. In the left pane, click Chassis Overview > Storage > Virtual Disks > Manage.
- 2. On the **Manage Virtual Disks** page, from the **Virtual Disk Actions** drop-down menu, select one of the following actions, and then click **Apply**.
 - Rename
 - Delete
 - NOTE: If you select **Delete**, the following message is shown indicating that deleting a virtual disk permanently deletes data available in that virtual disk.

Deleting the virtual disk removes the virtual disk from the controller's configuration. Initializing the virtual disk permanently erases data from the virtual disk.

Edit Policy: Read Cache
 Edit Policy: Write Cache
 Edit Policy: Disk Cache

Initialize: FastInitialize: FullEncrypt Virtual Disk

Enclosure Management Module

Enclosure Management Module (EMM) provides data path and enclosure management tasks for enclosure. EMM monitors and controls the enclosure components and access to the drives.

EMM communicates enclosure attributes and states to the host server. The EMM modules monitor the following components of the enclosure:

- Fans
- Power supplies

- Temperature probes
- Insertion or removal of a physical disk
- LEDs on the enclosure

Viewing EMM Status and attributes

EMM status displays the health of the EMM. EMMs contain a status value that is unique from the enclosure. You can have up to two EMMs. Enclosure firmware creates a status for each EMM.

Viewing EMM Status and Attributes Using Web Interface

To view the status and attributes of the EMM:

Click **Chassis Overview** \rightarrow **Storage** \rightarrow **Enclosures** \rightarrow **Properties**. The **Enclosures page** provides the EMM status and attributes of the enclosures in the chassis. Expand the integrated enclosure or external enclosures to view the status and attributes of the EMM. For more information, see the *CMC Online Help*.

Viewing EMM Status and Attributes Using RACADM

To view the status of EMM, use racadm raid get emms -o -p Status command.

To view the attributes of EMM, use racadm raid get emms -o command.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Viewing Enclosure Status and Attributes

CMC displays the health of the enclosure based on the physical components. Data of the enclosures attached to shared storage is displayed in CMC, but the external enclosures attached to few PCle cards are not displayed. You must have CMC Login privileges to view the status and attributes of the enclosures.

Viewing Enclosure Status and Attributes Using Web Interface

To view the status and attributes of the enclosure:

Click Chassis Overview \rightarrow Storage \rightarrow Enclosures \rightarrow Properties. The Enclosures page provides the health status of the enclosures in the chassis. For more information, see the CMC Online Help.

Viewing Status and Attributes of Enclosure Using RACADM

To view the status of the enclosure, use racadm raid get enclosures -o -p Status command.

To view the attributes of the enclosure, use racadm raid get enclosures —o command.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Reporting up to two Enclosures per Connector

Each External Shared PERC 8 card supports up to two enclosures per connector. However, there are two different configurations with different restrictions. In a single PERC (non-fault tolerant) configuration you can connect up to two enclosures per card. Because of redundant cabling, a Fault Tolerant External Shared PERC 8 card solution supports up to two enclosures per fault tolerant pair.

If more than two enclosures are detected on any connector, a warning message is logged to the chassis log. This affects chassis health and provides an active alert or chassis log entry.

Setting Asset Tag and Asset Name of the Enclosure

To identify the enclosures, set the asset name and asset tag of the enclosures.

(i) NOTE:

- Error is displayed if you enter an invalid value.
- Initially the value which is saved in the firmware is displayed.
- You must have Chassis Configuration Privileges to set asset tag and asset name of the enclosure.
- You can set Asset Tag and Asset Name only for external enclosures.

Setting Asset Tag and Asset Name of the Enclosure Using Web Interface

To set the asset tag and asset name of the enclosure, click **Chassis Overview** \rightarrow **Storage** \rightarrow **Enclosures** \rightarrow **Setup**. Type the **Asset Tag** and **Asset Name** in the appropriate fields, and then click **Apply**. For more information, see the *CMC Online Help*.

Setting Asset Tag and Asset Name of the Enclosure Using RACADM

To set asset tag of the enclosure, use racadm raid set enclosures: Enclosure.External.O-0:RAID.ChassisSlot.5-1 -p AssetTag <value> command.

To set asset name of the enclosure, use racadm raid set enclosures: Enclosure.External.0-0:RAID.ChassisSlot.5-1 -p AssetName <value> command.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Viewing Temperature Probe Status and attributes of the Enclosure

Temperature probe status displays the status of the temperature sensors of the enclosure. Sensors contain a status value that is unique from the enclosure. You can have up to four temperature sensors or probes and enclosure firmware creates a status for each sensor. You must have CMC Login privileges to view probe status.

Viewing Temperature Probe Status and attributes of the Enclosure Using Web Interface

To view the temperature probe status and attributes of the enclosure:

Click **Chassis Overview** \rightarrow **Storage** \rightarrow **Enclosures** \rightarrow **Properties**. The **Enclosures page** provides the health and attributes for the temperature probe of the enclosure in the chassis. Expand the external enclosure to view the status for the PSU of the enclosures. For more information, see the *CMC Online Help*.

Viewing Temperature Probe Attributes of the Enclosure Using RACADM

To view temperature probe attributes of the Enclosure, use racadm raid get tempprobes -o command. For more information, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at **dell.com/cmcmanuals**.

Setting the Temperature Warning Threshold of the Enclosure

Temperature warning threshold allows you to change the threshold at which an enclosure temperature reports as warning.

(i) NOTE:

- Error is displayed if you enter an invalid value.
- Initially the value which is saved in the firmware is displayed.
- You must have Chassis Configuration Privileges to set asset tag and asset name of the enclosure.

Setting the Temperature Warning Threshold of the Enclosure Using Web Interface

To set temperature warning threshold of the enclosure:

Click Chassis Overview \rightarrow Storage \rightarrow Enclosures \rightarrow Setup. Select the enclosure from the Enclosure drop-down menu, then enter the appropriate values for minimum and maximum for warning threshold temperatures of temp sensor 2 and 3. Type the Asset tag and Asset name in the appropriate fields, and then click Apply. For more information, see the CMC Online Help.

Setting Temperature Warning Threshold of the Enclosure Using RACADM

To set the minimum warning threshold of temperature probe in the enclosure, use racadm raid set tempprobes: TempSensor. Embedded. 0: Enclosure. External. 1-0: RAID. ChassisSlot. 6-1 -p MinimumWarningThreshold <value> command.

To set the maximum warning threshold of temperature probe in the enclosure, use racadm raid set tempprobes: TempSensor. Embedded. 0: Enclosure. External. 1-0: RAID. ChassisSlot. 6-1 -p MaximumWarningThreshold <value> command.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Viewing Fan Status and attributes of the Enclosure

Fan status and attributes displays the status of the enclosure fan and contains a status value that is unique from the enclosure. You can have up to two fans, and the enclosure firmware creates a status for each fan. You must have CMC Login privileges to view fan status.

i NOTE: If a PSU is missing, the corresponding fan of the PSU displays a critical status.

Viewing Fan Status and attributes of the Enclosure Using Web Interface

To view the status and attributes of PSU:

Click Chassis Overview \rightarrow Storage \rightarrow Enclosures \rightarrow Properties. The Enclosure Page provides the health status and attributes for the fan of the enclosure. Expand the external enclosure to view the status for the fan of the enclosure. For more information, see the CMC Online Help.

Viewing Fan Status and attributes of the Enclosure Using RACADM

To view the status of fan, use racadm raid get fans -o -p Status command.

To view the attributes of fan, use ${\tt racadm}$ ${\tt raid}$ get fans ${\tt -o}$ command.

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide available at dell.com/support/manuals.

Viewing Enclosure Properties Using CMC Web Interface

To view the enclosure properties:

- 1. In the left pane, click Chassis Overview > Storage > Enclosures > Properties.
- 2. On the **Properties** page, under the **Enclosure** section, click the |+| to get a graphical view of the physical disk drives and their states, summary of physical disk drive slots, and advanced properties.

Managing PCIe Slots

By default, all slots are unmapped. You can do the following:

- View the status of all PCle Slots in the chassis.
- Assign or remove an assigned PCle slot from the servers.

Consider the following before assigning a PCIe slot to a server:

- An empty PCle slot cannot be assigned to a server that is turned on.
- A PCle slot with an adapter assigned to a server cannot be assigned to another server if the currently-assigned server (source) is turned on.
- A PCIe slot with an adapter assigned to a server cannot be assigned to another (target) server which is turned on.

Consider the following before removing an assigned PCle slot from a server:

- If a PCle slot is empty, a slot can be unassigned from a server, even if the server is turned on.
- If a PCle slot has adapter and it is not turned on, it can be unassigned from the server even if the server is on. This can occur when a slot is empty and the assigned server is turned on, and then a user inserts adapter in an empty slot.

Map or un-map the external PCle adapter to blade:

- Adapter is always powered on as a non-shared device. Henceforth, adapter is mapped to one server.
- If an external PCIE slot is occupied by a shared adapter, the mapping prior to adapter which is inserted remains unchanged.
- If an external PCIE slot is occupied by a shared adapter, PCIE slot may not map or unmap to or from a blade server. If a user attempts to map or unmap shared adapter an EEMI message is logged.

For more information about assigning and remove an assigned PCle slot from the servers, see the Online Help.

(i) NOTE:

- Without a license, you can assign a maximum of four PCle slots to full-height server, two to the upper slot and two to the extension slot, or two PCle devices to a half-height server.
- You can distinguish the properties of external PCIE slots with External Shared PERC 8 card devices from dedicated devices, these shared devices have different properties than dedicated devices.
- In case of external SPERC, the status Shared is displayed. The options to map or unmap the External Shared PERC 8 card is not available.

Topics:

- Viewing PCle Slot Properties Using CMC Web Interface
- Assigning PCle Slots To Servers Using CMC Web Interface
- Managing PCle Slots Using RACADM
- PCle Power Ride-Through

Viewing PCIe Slot Properties Using CMC Web Interface

- To view the information about all the eight PCle slots, in the left pane, click **Chassis Overview** > **PCle Overview**. Click the to view all the properties for the required slot.
- To view the information about one PCle slot, click Chassis Overview > PCle Slot <number> > Properties > Status.
- NOTE: User Interface differentiates the external PCIe slots which contains SPERC (or any shared) devices installed from external PCIE slots with dedicated adapters as these shared devices have different properties.

Assigning PCIe Slots To Servers Using CMC Web Interface

To assign PCle slots to the servers:

In the left pane, click Chassis Overview > PCle Overview > Setup > Mapping: PCle Slots to Server Slots. On
the Mapping: PCle Slots to Server Slots page, in the Action column, from the Action drop-down menu, select the
appropriate server name, and then click Apply.

Note the following:

- Without a license, the maximum number of PCle slots that maybe mapped to a half-height server is two. If a full-height server is installed you can map two PCle slots to the upper server slot and two to the lower (extended) server slot, for a total of four PCle slots per full-height server.
- You can map the server slots to any of the 8 PCle slots.
- A full-height server has both upper and lower mezzanines populated. Else, during POST will stop when the <F1> or <F2> is displayed on the page for you to press any one of the keys.
- For full-height servers you can map a maximum of two PCIE slots to upper and two to lower mezzanines. By default, all PCIe mappings to server slot 3 will go to the lower mezzanines.
- Server slot number is displayed as Slot-01, Slot-02, and so on. For a full-height server, the slot name is displayed as Ext. of Slot-01, Ext. of Slot-02, and so on.
- If you select the host name, then the host name is displayed instead of the slot name.
- CMC provides alert capabilities through System Event Log (SEL), SNMP, and Email interfaces.

For more information about assigning PCle devices to a server, see the Online Help.

Managing PCIe Slots Using RACADM

You can assign or unassign a PCle slot to a server by using the RACADM commands. Some of the commands are given here. For more information about the RACADM commands, see the *Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide* available at dell.com/support/Manuals.

• To view the current assignment of PCle devices to servers, run the following command:

```
racadm getpiecfg -a
```

• To view the properties of PCIe devices by using FQDD, run the following command:

```
racadm getpciecfg [-c <FQDD>]
```

For example, to view the properties of PCIe device 1, run the following command.

```
racadm getpciecfg -c PCIE.ChassisSlot.1
```

To assign a PCle adapter slot to server slot, run the following command:

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

• For example, to assign PCle slot 5 to server slot 2, run the following command.

```
racadm setpciecfg assign -c PCIE.ChassisSlot.5 -i 2
```

• To unassign a PCle slot 3 from a server, run the following command:

```
racadm setpciecfg unassign -c pcie.chassisslot.3
```

PCIe Power Ride-Through

Newly assigned PCIe cards in CMC VRTX must be discovered and initialized before a server node is powered on. The discovery and initialization process involves the following:

- Inventory and discovery of installed cards
- Preparing a PCle card for exposure to a server module
- Preparing multiple cards for configuration by server BIOS
- Initializing all cards prior to blade node power-on

All these processes require few seconds to complete which causes a delay in initialization of the PCle cards. The PCle Ride-through feature in CMC VRTX reduces this process cycle time. The PCle Ride-through feature enables the following:

- The Server nodes are turned on quickly, thus turning on the PCle cards quickly.
- The powered state of PCIe cards is extended for a pre-defined time period in the following scenarios:
 - o After the associated server is turned off.
 - After adapter discovery process is completed.
- The power-on readiness state of the cards is extended for a predefined time after the discovery process. This extension
 eliminates the delays for common types of power cycling scenarios. The cards continue to be in ready state awaiting node
 assignment and power-up. The cards power-down after the time period expires.
- **NOTE:** At the end of the time period, the PCIe cards power-down. All adapters in ride-through are also powered down whenever the chassis door is opened.
- (i) NOTE:
 - If CMC has insufficient power, the CMC turns off all adapters in ride-through mode, thereby releasing all power allocated to those adapters. If the power supply is restored, the power is reallocated to the PCle slots. This power restoration enables the cards to be ready for server allocation without delay.
 - All external PCIE adapters powered on as shared mode is excluded from the ride-through processes. After a shared adapter is powered on as a shared device, it remains powered on until the chassis is powered off.

Viewing PCIe Ride-through Properties Using CMC Web Interface

To view the PCle Ride-through properties, in the left pane, click **Chassis Overview** > **PCle Overview**. The **PCle Status** page is displayed. The **General Settings** section displays the following PCle Ride-through properties status:

- Ridethrough Status Enabled or Disabled
- Ridethrough Timeout Indicates the time for which the Ride-through feature is enabled

Viewing PCIe Ridethrough Properties Status Using RACADM

To view the information about PCle power ridethrough properties, enter the following command:

racadm getpciecfg -r

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Configuring PCIe Ride-through Properties Using CMC Web Interface

To configure the PCle Ride-through properties for CMC VRTX:

In the left pane, click Chassis Overview > Setup > Ridethrough.
 The PCIe Ridethrough Settings page is displayed.

- 2. To enable or disable the PCle Ride-through feature, select or clear the Enable PCle Ridethrough option.
 - (i) NOTE: By default, the Ride-through feature is enabled and the time period set for 300 seconds.
- 3. In the **Timeout** field, enter the time for which the Ride-through feature to be enabled. Type either zero (0) or a value from 60–1800 seconds. Zero indicates an infinite timeout.
- 4. Click Apply.

Configuring PCIe Ride-through Properties Status Using RACADM

You can configure the PCle power ride-through properties, by running the following commands:

- To disable the Ride-through feature, run the command, racadm setpciecfg ridethru -d
- To enable the Ride-through feature, run the command, racadm setpciecfg ridethru -e
- To reset the Ride-through timeout property, run the command, racadm setpciecfg ridethru -t <timeout>
- To set the acceptable timeout range, run the command, racadm setpciecfg help ridethru

For more information, see the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

Troubleshooting and Recovery

This section explains how to perform tasks related to recovering and troubleshooting problems on the remote system using the CMC web interface.

- Viewing chassis information.
- Viewing the event logs.
- Gathering configuration information, error status, and error logs.
- Using the diagnostic console.
- Managing power on a remote system.
- Managing Lifecycle Controller jobs on a remote system.
- Reset components.
- Troubleshooting Network Time Protocol (NTP) problems.
- Troubleshooting network problems.
- Troubleshooting alerting problems.
- Resetting forgotten administrator password.
- Saving and restoring Chassis configuration settings and certificates.
- Viewing error codes and logs.

i NOTE: WinRM support for Microsoft is not available for Windows 10 client, use Power Shell instead of WinRM.

Topics:

- Resetting Forgotten Administrative Password
- Gathering Configuration Information, Chassis Status, and Logs Using RACDUMP
- First Steps to Troubleshoot a Remote System
- Troubleshooting Alerts
- Viewing Event Logs
- Using Diagnostic Console
- Resetting Components
- Saving or Restoring Chassis Configuration
- Troubleshooting Network Time Protocol Errors
- Interpreting LED Colors and Blinking Patterns
- Troubleshooting Non-responsive CMC
- Troubleshooting Network Problems
- Troubleshooting Controller
- Hotplugging enclosures in fault-tolerant chassis

Resetting Forgotten Administrative Password

The following procedure explains how to reset the forgotten administrative password:

- Remove CMC module from the chassis
- Short the Password recovery header pins using jumper
- Re-insert the CMC module to chassis, after the CMC is in online state, the default credential is active (User Name: root/ Password: calvin)
- Login to CMC using default credential and change the password
- After the password is changed, remove CMC module and jumper from Password recovery header
- Re-insert the CMC module to chassis, after the CMC is in online state, the new credential is active

Gathering Configuration Information, Chassis Status, and Logs Using 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 following information:

- General system/RAC information
- CMC information
- Chassis information
- Session information
- Sensor information
- Firmware build information

Supported Interfaces

- CLI RACADM
- Remote RACADM
- Telnet RACADM

racdump includes the following subsystems and aggregates the following RACADM commands. For more information about racdump, see the RACADM Command Line Reference Guide for CMC in PowerEdge VRTX.

Table 41. Racadm commands for subsystems

Subsystem	RACADM Command
General System/RAC information	getsysinfo
Session information	getssninfo
Sensor information	getsensorinfo
Switches information (IO Module)	getioinfo
Mezzanine card information (Daughter card)	getdcinfo
All modules information	getmodinfo
Power budget information	getpbinfo
KVM information	getkvminfo
NIC information (CMC module)	getniccfg
Redundancy information	getredundancymode
Trace log information	gettracelog
RAC event log	getraclog
System event log	getsel

Downloading SNMP Management Information Base File

The CMC SNMP Management Information Base (MIB) file defines the chassis types, events, and indicators. CMC enables you to download the MIB file using the web interface.

To download the CMC's SNMP MIB file using the CMC web interface:

- 1. In the left pane, click Chassis Overview > Network > Services > SNMP.
- 2. In the **SNMP Configuration** section, click **Save** to download the CMC MIB file to your local system.

 For more information about the SNMP MIB file, see the *Dell OpenManage Server Administrator SNMP Reference Guide* at **dell.com/support/manuals.**

First Steps to Troubleshoot a Remote System

The following questions are commonly used to troubleshoot high-level issues in the managed system:

- Is the system turned on or turned off?
- If turned on, is the operating system functioning, not responding, or stopped functioning?
- If turned off, did the power turn off unexpectedly?

Power Troubleshooting

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

- **Problem:** Configured the **Power Redundancy Policy** to **Grid Redundancy**, and a Power Supply Redundancy Lost event was raised.
 - **Resolution A:** This configuration requires at least one power supply in side 1 (the left two slots) and one power supply in side 2 (the right two slots) to be present and functional in the modular enclosure. Additionally the capacity of each side must be enough to support the total power allocations for the chassis to maintain **Grid redundancy**. (For full Grid Redundancy operation, make sure that a full PSU configuration of four power supplies is available.)
 - Resolution B: Check if all power supplies are properly connected to the two AC grids; power supplies in side 1 must be connected to one AC grid, those in side 2 must be connected to the other AC grid, and both AC grids must be working.
 Grid 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.
 - **Resolution A:** Check and replace the AC cord. Check and confirm that the power distribution unit providing power to the power supply is operating as expected. If the failure still persists, call Dell customer service for replacement of the power supply.
 - **Resolution B:** Check that the PSU is connected to the same voltage as the other PSUs. If CMC detects a PSU operating at a different voltage, the PSU is turned off and marked Failed.
- Problem: Dynamic Power Supply Engagement is enabled, but none of the power supplies display in the Standby state.
 - **Resolution A:** There is insufficient surplus power. One or more power supplies are moved into the Standby state only when the surplus power available in the enclosure exceeds the capacity of at least one power supply.
 - **Resolution B:** Dynamic Power Supply Engagement cannot be fully supported with the power supply units present in the enclosure. To check if this is the case, use the web interface to turn Dynamic Power Supply Engagement off, and then on again. A message is displayed if Dynamic Power Supply Engagement cannot be fully supported.
- Problem: Inserted a new server into the enclosure with sufficient power supplies, but the server does not power on.
 - **Resolution A:** Check for the system input power cap setting—it might be configured too low to allow any additional servers to be powered up.
 - **Resolution B:** Check for the maximum power conservation setting. If this is set, then this issue occurs. For more details, see the power configuration settings.
 - **Resolution C:** Check for the server slot power priority of the slot associated with the newly-inserted server, and then ensure it is not lesser than any other server slot power priority.
- Problem: Available power keeps changing, even when the modular enclosure configuration has not changed.
 - **Resolution:** CMC has dynamic fan power management that reduces server allocations briefly if the enclosure is operating near the peak user configured power cap; it causes the fans to be allocated power by reducing server performance to keep the input power draw below **System Input Power Cap**. This is normal behavior.
- Problem: <number>W is reported as the Surplus for Peak Performance.
 - **Resolution:** The enclosure has <number>W of surplus power available in the current configuration, and the **System Input Power Cap** can be safely reduced by this amount being reported without impacting server performance.
- **Problem:** A subset of servers lost power after an AC Grid failure, even when the chassis was operating in the **Grid Redundancy** configuration with four power supplies.
 - Resolution: This can occur if the power supplies are improperly connected to the redundant AC grids at the time the AC grid failure occurs. The **Grid Redundancy** policy requires that the left two power supplies be connected to one AC grid, and right two power supplies be connected to other AC grid. If two PSUs are improperly connected, such as PSU 2 and PSU 3 are connected to the wrong AC grids, an AC grid failure cause loss of power to the least priority servers.
- **Problem:** The least priority servers lost power after a PSU failure.
 - Resolution: To avoid a future power supply failure causing servers to power off, make sure that the chassis has at least
 three power supplies 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 has been configured to a value that results in an increased power need 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.

Troubleshooting Alerts

Use the CMC log and the trace log to troubleshoot CMC alerts. The success or failure of each email and/or SNMP trap delivery attempt is logged into the CMC log. Additional information describing the particular error is logged in the trace log. However, since SNMP does not confirm delivery of traps, use a network analyzer or a tool such as Microsoft's snmputil to trace the packets on the managed system.

Viewing Event Logs

You can view hardware- and chassis logs for information on system-critical events that occur on the managed system.

Viewing Hardware Log

CMC generates a hardware log of events that occur on the chassis. You can view the hardware log using the web interface and remote RACADM.

- i NOTE: To clear the hardware log, you must have Clear Logs Administrator privilege.
- NOTE: You can configure CMC to send email or SNMP traps when specific events occur. For information on configuring CMC to send alerts, see Configuring CMC To Send Alerts.

Examples of hardware log entries

```
critical System Software event: redundancy lost
Wed May 09 15:26:28 2007 normal System Software
event: log cleared was asserted
Wed May 09 16:06:00 2007 warning System Software
event: predictive failure was asserted
Wed May 09 15:26:31 2007 critical System Software
event: log full was asserted
Wed May 09 15:47:23 2007 unknown System Software
event: unknown event
```

Viewing Hardware Logs Using CMC Web Interface

You can view, save, and clear the hardware log. You can sort the log entries based on Severity, Date/Time, or Description by clicking the column heading. Subsequent clicks on the column headings reverse the sort.

To view the hardware logs using CMC Web interface, in the left pane, click **Chassis Overview** > **Logs**. The **Hardware Log** page is displayed. To save a copy of the hardware log to your managed station or network, click **Save Log**, and then specify a location for a text file of the log.

NOTE: Since the log is saved as a text file, the graphical images used to indicate severity in the user interface do not appear. In the text file, severity is indicated with the words **OK**, **Informational**, **Unknown**, **Warning**, and **Severe**. The date and time entries appear in ascending order. If <SYSTEM BOOT> appears in the **Date/Time** column, it means that the event occurred during the turn-on or turn—off of any of the modules, when a date or time is not available.

To clear the hardware log, click Clear Log.

- i NOTE: CMC creates a new log entry indicating that the log was cleared.
- i NOTE: To clear the hardware log, you must have the Clear Logs Administrator privilege.

Viewing Hardware Logs Using RACADM

To view the hardware log using RACADM, open a serial/Telnet/SSH text console to CMC, log in, and type:

racadm getsel

To clear the hardware log, type:

racadm clrsel

Viewing Chassis Log

CMC generates a log of the chassis-related events. CMC provides alert capabilities through System Event Log (SEL), SNMP, and Email interfaces.

SPERC is inserted while one or more PowerEdge serves are powered on.

(i) NOTE:

• To clear the chassis log, you must have the Clear Logs Administrator privilege.

Viewing Chassis Logs Using RACADM

To view the chassis log information using RACADM, open a serial, Telnet, SSH text console to CMC, log in, and enter the following:

racadm chassislog view

This command displays the latest 25 chassis log entries.

To display the options available to view chassislogs, run the following command:

racadm chassislog help view

Viewing Chassis Logs Using the Web Interface

You can view, save, and clear the chassis log. You can filter the logs based on the log type and filter. Additionally, you can even perform a search based on a keyword or view the logs on specified days.

In the left pane, click Chassis Overview > Logs > Chassis Log. The Chassis Log page is displayed.

To save a copy of the chassis log to your managed station or network, click **Save Log** and then specify a location save the log file

Using Diagnostic Console

You can diagnose issues related to the chassis hardware using CLI commands if you are an advanced user or a user under the direction of technical support.

i NOTE: To modify these settings, you must have the **Debug Command Administrator** privilege.

To access the Diagnostic Console:

- In the left pane, click Chassis Overview > Troubleshooting > Diagnostics.
 The Diagnostic Console page displays.
- 2. In the Command text box, type a command and click Submit.

For information about the commands, see the Online Help.

The diagnostic results page appears.

Resetting Components

You can reset the active CMC, or virtually reseat servers making 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.

i NOTE: To reset components, you must have Debug Command Administrator privilege.

To reset the components using the CMC Web interface,

- In the left pane, click Chassis Overview > Troubleshooting > Reset Components.
 The Reset Components page is displayed.
- 2. 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. However, if a standby CMC is not present, the CMC that is available is rebooted.
- To virtually reseat the server, in the Virtual Reseat Server section, select servers to reseat, and then click Apply Selections.

For more information, see the Online Help.

This operation causes the servers to behave as if they were removed and reinserted.

Saving or Restoring Chassis Configuration

This is a licensed feature. To save or restore a backup of the Chassis configuration using the CMC Web interface:

- 1. In the left pane, click **Chassis Overview** > **Setup** > **Chassis Backup**. The **Chassis Backup** page is displayed. To save the chassis configuration, 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.
- 2. To restore the chassis configuration, in the "Restore" section, click Browse, specify the backup file, and then click Restore.

(i) NOTE:

- CMC does not reset upon restoring configuration, however CMC services may take some time to effectively impose any changed or new configuration. After successful completion, all current sessions are closed.
- Flexaddress information, server profiles, and extended storage are not saved or restored with the Chassis Configuration.

Troubleshooting Network Time Protocol Errors

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

- Problem with the Network Time Protocol (NTP) Server 1, NTP Server 2, and NTP Server 3 settings.
- Invalid host name or IP address may have been accidentally entered.
- Network connectivity problem 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 the NTP-related problems, check the information in the CMC trace log. This log contains an error message for NTP related failures. 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 typing the following racadm command:

```
racadm getractime -n
```

The output of this command contains detailed NTP statistics that may be useful in debugging the problem.

If you attempt 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, since the default setting must be large enough to work with most NTP servers.

To modify the parameter, type the following command:

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

After making the change, disable NTP, wait for 5-10 seconds, then enable NTP again:

(i) NOTE: NTP may take an additional three minutes to synchronize again.

To disable NTP, type:

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

To enable NTP, type:

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

For information on how to enter the gettracelog command to review the trace log using the CMC Web interface, see Using Diagnostic Console.

Interpreting LED Colors and Blinking Patterns

The LEDs on the chassis provide the following status of a component:

- 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. For more information about configuration, see Configuring LEDs to Identify Components on the Chassis.

Table 42. LED Color and Blinking Patterns

Component	LED Color, Blinking Pattern	Status
CMC	Green, glowing steadily	Turned on
	Green, blinking	Firmware is being uploaded
	Green, dark	Turned off
	Blue, glowing steadily	Active
	Blue, blinking	User-enabled module identifier
	Amber, glowing steadily	Not used
	Amber, blinking	Fault
	Blue, dark	Standby
Server	Green, glowing steadily	Turned on

Table 42. LED Color and Blinking Patterns (continued)

Component	LED Color, Blinking Pattern	Status
	Green, blinking	Firmware is being uploaded
	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	Turned on
	Green, blinking	Firmware is being uploaded
	Green, dark	Turned off
	Blue, glowing steadily	Normal/stack master
	Blue, blinking	User-enabled module identifier
	Amber, glowing steadily	Not used
	Amber, blinking	Fault
	Blue, dark	No fault/stack slave
IOM (Pass through)	Green, glowing steadily	Turned on
	Green, blinking	Not used
	Green, dark	Powered off
	Blue, glowing steadily	Normal
	Blue, blinking	User-enabled module identifier
	Amber, glowing steadily	Not used
	Amber, blinking	Fault
	Blue, dark	No fault
Blower	Green, glowing steadily	Fan working
	Green, blinking	Not used
	Green, dark	Turned off
	Amber, glowing steadily	Fan type not recognized, update the CMC firmware
	Amber, blinking	Fan fault; tachometer out of range
	Amber, dark	Not used
PSU	(Oval) Green, glowing steadily	AC OK
	(Oval) Green, blinking	Not used
	(Oval) Green, dark	AC Not OK
	Amber, glowing steadily	Not used
	Amber, blinking	Fault
	Amber, dark	No fault
	(Circle) Green, glowing steadily	DC OK
	(Circle) Green, dark	DC Not OK

Table 42. LED Color and Blinking Patterns (continued)

Component	LED Color, Blinking Pattern	Status
Enclosure	Blue	When the host server is identifying the enclosure
	Amber	Turned on or Reset, fault state

Troubleshooting Non-responsive CMC

If you cannot log in to CMC using any of the interfaces (the web interface, Telnet, SSH, remote RACADM, or serial), you can verify the CMC functionality observing the LEDs on CMC, obtaining recovery information using the DB-9 serial port, or recovering the CMC firmware image.

i NOTE: It is not possible to log in to the standby CMC using a serial console.

Observing LEDs to Isolate the Problem

There are two LEDs on the left of the card:

- Upper-left LED Indicates power status. If it is not ON:
 - o Verify that you have AC present to at least one power supply.
 - Verify that the CMC card is seated properly. You can release or pull the ejector handle, remove the CMC, reinstall the CMC making sure that the board is inserted all the way and the latch closes correctly.
- Lower-left LED This 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 is detected. The fault may be caused by any of the following three events:
 - o A core failure. In this case, the CMC board must be replaced.
 - o A self-test failure. In this case, the CMC board must be replaced.
 - o An image corruption. In this case, upload the CMC firmware image to recover the CMC.
 - NOTE: A normal CMC boot or reset takes over a minute to fully boot into its operating system and be available for login. The blue LED is enabled on the active CMC. In a redundant, two-CMC configuration, only the upper-right green LED is enabled on the standby CMC.

Obtain Recovery Information from DB-9 Serial Port

If the bottom LED is amber, recovery information is available from the DB-9 serial port located on the front of CMC.

To obtain recovery information:

- 1. Install a NULL modem cable between a CMC system and a client system.
- 2. Open a terminal emulator of your choice (such as HyperTerminal or Minicom). Enter the following specification when prompted: 8 bits, no parity, no flow control, baud rate 115200.
- **3.** Press the <Enter> key.

If a recovery prompt appears, additional information is available. The prompt indicates the CMC slot number and failure type. To display failure reason and syntax for a few commands, type recover, and then press <Enter>.

Sample prompts:

```
recover1[self test] CMC 1 self test failure
recover2[Bad FW images] CMC2 has corrupted images
```

- If the prompt indicates a self test failure, there are no serviceable components on CMC. CMC is bad and must be returned to Dell.
- If the prompt indicates Bad FW Images, complete tasks in Recovering Firmware Image.

Recovering Firmware Image

CMC enters recover mode when a normal CMC operating boot is not possible. 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 and then press <Enter> at the recovery prompt, the recover reason and available sub-commands display. An example recover sequence may be:

```
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
```

- (i) NOTE: Connect the network cable to the left most RJ45.
- NOTE: 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 may need to use the recover reset command after setniccfg on some systems.

Troubleshooting Network Problems

The integrated CMC trace log allows you to debug CMC alerts and networking. You can access the trace log using the CMC Web interface or RACADM. See the <code>gettracelog</code> command section in the Chassis Management Controller for PowerEdge VRTX RACADM Command Line Reference Guide.

The trace log tracks the following information:

- DHCP Traces packets sent to and received from a DHCP server.
- DDNS Traces dynamic DNS update requests and responses.
- Configuration changes to the network interfaces.

The trace log may also contain CMC firmware-specific error codes that are related to the integrated CMC firmware, not the managed system's operating system.

Troubleshooting Controller

To troubleshoot a controller:

- 1. In the left pane, click Chassis Overview > Storage > Controllers > Troubleshooting.
- 2. On the **Controller Troubleshooting** page, from the **Actions** drop-down list for the respective controller, select any one of the following, and then click **Apply**.
 - Reset Configuration Deletes the virtual disks and hot spares. However, the data on the disks is not erased.
 - (i) NOTE: Resetting PERC configuration discards the pinned cache, if any, on the PERC controller.
 - Export TTY Log The TTY debug log from the storage controller is exported to your local system.
 - **Discard Pinned Cache** Deletes data that is stored in RAID controller cache.
 - NOTE: If there is pinned cache, the option to clear it is present. If there is no pinned cache, this option is not displayed.
 - **Disable RAID Controller** Disables the peer controller. This option is available in the drop-down menu only for the Shared PERC8 (Integrated 2) and external Shared PERC8s.
 - Enable RAID Controller Enables the peer controller. Enable Raid Controller option is available in the drop-down menu.
 - (i) NOTE:

For a disabled PERC, none of the other options Reset Configuration, Export TTY Log, Discard Pinned Cache, and Disable RAID Controller are available in the drop-down menu.

• Enable Fault Tolerance — Enables the fault-tolerance mode of the External shared PERC 8 card.

- **Disable Fault Tolerance** Disables the fault-tolerance mode of the External shared PERC 8 card.
 - NOTE: Enable Fault Tolerance and Disable Fault Tolerance are displayed only for the External Shared PERC 8 cards. The default mode of the External Shared PERC 8 cards is the non-fault-tolerant mode.

(i) NOTE:

- Displays an error message if the blades are powered ON.
- The command fails if the blade is powered ON.

Hotplugging enclosures in fault-tolerant chassis

- 1. Make sure that slots 5 and 6 chassis are not fault-tolerant.
- 2. Disconnect the enclosures.
- 3. Change the status of slots 5 and 6 to fault-tolerant mode.
- **4.** Reconnect the enclosures in fault-tolerant cabling.

Power cycle the chassis after disconnecting the enclosures and before reconnecting enclosures as, the drives retain the previous SCSI-3 Reservation until the chassis is power cycled.

Using LCD Panel Interface

You can use the LCD panel on the chassis to perform configuration and diagnostics, and to obtain status information about the chassis and its contents.

The following figure illustrates the LCD panel. The LCD screen displays menus, icons, pictures, and messages.




Figure 4. LCD Display

- 1. LCD screen
- 2. Selection ("check") button
- 3. Scroll buttons (4)

Topics:

- LCD Navigation
- Diagnostics
- Front Panel LCD Messages
- LCD Module and Server Status Information

LCD Navigation

The right side of the LCD panel contains five buttons: four arrow buttons (up, down, left, and right) and a center button.

- To move between screens, use the right (next) and left (previous) arrow buttons. At any time while using the panel, you can return to a previous screen.
- To scroll through options on a screen, use the down and up arrow buttons.
- To select and save an item on a screen and move to the next screen, use the center button.

The up, down, left, and right arrow buttons change the selected menu items or icons on the screen. The selected item is shown with a light blue background or border.

When messages displayed on the LCD screen are longer than what fits on the screen, use the left and right arrow buttons to scroll the text left and right.

The icons described in the following table are used to navigate between LCD screens.

Table 43. LCD Panel Navigational Icons

Icon Normal	lcon Highlighted	Icon Name and Description
•		Back — Highlight and press the center button to return to the previous screen.

Table 43. LCD Panel Navigational Icons (continued)

\checkmark		Accept/Yes — Highlight and press the center button to accept a change and return to the previous screen.
		Skip/Next — Highlight and press the center button to skip any changes and go to the next screen.
$\overline{\mathbf{X}}$		No — Highlight and press the center button to answer "No" to a question and go to the next screen.
	***	Component Identify — Blinks the blue LED on a component. i NOTE: There is a blinking blue rectangle around this icon when Component Identify is enabled.

A status indicator LED on the LCD panel provides an indication of the overall health of the chassis and its components.

- Solid blue indicates good health.
- Blinking amber indicates that at least one component has a fault condition.
- Blinking blue is an ID signal, used to identify one chassis in a group of chassis.

Main Menu

From the Main Menu menu, you can navigate to one of the following screens:

- KVM Mapping Contains the options to map or unmap the KVM to the servers.
- DVD Mapping This option is displayed on the Main Menu screen only if you a DVD drive installed.
- **Enclosure** Displays status information for the chassis.
- IP Summary Displays information about CMC IPv4, CMC IPv6, iDRAC IPv4, and iDRAC 4 IPv6.
- Settings Contains the options such as LCD Language, Chassis Orientation, Default LCD Screen, and the Network Settings.

KVM Mapping Menu

From this screen, you can view the KVM to server mapping information, map another server to the KVM, or unmap the existing connection. To use the KVM for a server, select **KVM mapping** from the main menu, navigate to the appropriate server, and then press the center **Check** button.

DVD Mapping

By using this page, you can view the DVD to server mapping information, map another server to the DVD drive on the chassis, or unmap the existing connection. To give a server access to the DVD, select **DVD mapping** from the main menu, navigate to the required server, and then press the center **Check** button.

The DVD drive can be mapped to the server slot only if the DVD is enabled for that server slot. DVD drive can also be unmapped to prevent the use by any of the server slots. The health of the DVD drive will be critical if the SATA cable is not properly connected between the DVD drive and the mainboard. If the health of the DVD drive is critical, the server cannot access the DVD drive.

i NOTE: The DVD Mapping feature is displayed on the LCD Main Menu screen only if you have a DVD drive installed.

Enclosure Menu

From this screen, you can navigate to the following screens:

- Front Status
- Rear Status

- Side Status
- Enclosure Status

Use the navigation buttons to highlight the desired item (highlight the **Back** icon to return to the **Main Menu**), and then press the center button. The selected screen is displayed.

IP Summary Menu

The **IP Summary** screen displays the IP information about the CMC (IPv4 and IPv6), and each server that is installed on the chassis.

Use the up and down arrow buttons to scroll through the list. Use the left and right arrow buttons to scroll selected messages that are longer than the screen.

Use the up and down arrow buttons to select the **Back** icon and press the center button to return to the **Enclosure** menu.

Settings

The **Settings** menu displays a menu of items that can be configured:

- LCD Language Select the language you want to use for LCD screen text and messages.
- Chassis Orientation Select either Tower Mode or Rack Mode on the basis of installation orientation of the chassis.
- Default LCD Screen Select the screen (Main Menu, Front Status, Rear Status, Side Status, or Custom) that is displayed when there is no activity on the LCD panel.
- Network Settings Select to configure the network settings of a CMC. For more information about this feature, see Configuring CMC Network Using LCD Panel Interface.

Use the up and down arrow buttons to highlight an item in the menu, or highlight the **Back** icon if you want to return to the **Main Menu** screen.

To activate your selection, press the center button.

LCD Language

The **LCD Language** screen allows you to select the language used for LCD panel messages. The currently active language is highlighted with a light blue background.

- 1. Use the up, down, left, and right arrow buttons to highlight the desired language.
- 2. Press the center button. The Accept icon appears and is highlighted.
- 3. Press the center button to confirm the change. The LCD Setup menu is displayed.

Default Screen

The **Default Screen** allows you to change the screen that the LCD panel displays when there is no activity at the panel. The factory default screen is the **Main Menu.** You can choose from the following screens to display:

- Main Menu
- Front Status (front graphical view of the chassis)
- Rear Status (rear graphical view of the chassis)
- Side Status (left graphical view of the chassis)
- Custom (Dell logo with chassis name)

The currently active default screen is highlighted in light blue color.

- 1. Use the up and down arrow keys to highlight the screen you want to set to the default.
- 2. Press the center button. The Accept icon is highlighted.
- 3. Press the center button again to confirm the change. The **Default Screen** is displayed.

Diagnostics

The LCD panel helps you to diagnose issues in any server or module in the chassis. If there is an issue or fault with the chassis or any server or other module in the chassis, the LCD panel status indicator blinks amber. On the **Main Menu**, an icon with an amber background displays next to the menu item— Enclosure—that leads to the Front, Rear, Side, or Enclosure status.

By following the amber icons through the LCD menu system, you can display the status screen and error messages for the item that has the issue.

Error messages on the LCD panel can be removed by removing the module or server that is the cause of the issue, or by clearing the hardware log for the module or server. For server errors, use the iDRAC web interface or command line interface to clear the server's System Event Log (SEL). For chassis errors, use the CMC web interface or command line interface to clear the hardware log.

Front Panel LCD Messages

This section contains two subsections that list error and status information that is displayed on the front panel LCD.

Error messages on the LCD have a format that is similar to the System Event Log (SEL) viewed from the CLI or Web interface.

The tables in the error section list the error and warning messages that are displayed on the various LCD screens and the possible cause of the message. Text enclosed in angled brackets (< >) indicates that the text may vary.

Status information on the LCD includes descriptive information about the modules in the chassis. The tables in this section describe the information that is displayed for each component.

LCD Module and Server Status Information

The tables in this section describe status items that are displayed on the front panel LCD for each type of component in the chassis.

Table 44. CMC Status

Item	Description
Name/Location	Example: CMC1, CMC2.
No Errors	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Firmware Version	Only displays on an active CMC. Displays Standby for the standby CMC.
IP4 <enabled, disabled=""></enabled,>	Displays current IPv4 enabled state only on an active CMC.
IP4 Address: <address, acquiring=""></address,>	Only displays if IPv4 is enabled only on an active CMC.
IP6 <enabled, disabled=""></enabled,>	Displays current IPv6 enabled state only on an active CMC.
IP6 Local Address: <address></address>	Only displays if IPv6 is enabled only on an active CMC.
MAC: <address></address>	Displays the CMC's MAC address.

Table 45. Chassis or Enclosure Status

Item	Description
User Define Name	Example: "Dell Rack System". This can be configured using the CMC CLI or Web interface.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Model Number	Example "PowerEdgeM1000".
Power Consumption	Current power consumption in watts.
Peak Power	Peak power consumed in watts.

Table 45. Chassis or Enclosure Status (continued)

Item	Description
Minimum Power	Minimum power consumed in watts.
Ambient Temperature	Current ambient temperature in degrees Celsius.
Service Tag	The factory-assigned service tag.
CMC redundancy mode	Non-Redundant or Redundant.
PSU redundancy mode	Non-Redundant, Grid Redundant, or DC Redundant.

Table 46. Fan Status

Item	Description
Name/Location	Example: Fan1, Fan2, and so on.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
RPM	Current fan speed in RPM.

Table 47. PSU Status

Item	Description
Name/Location	Example: PSU1, PSU2, and so on.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Status	Offline, Online, or Standby — Indicates the power status of a PSU.
Maximum Wattage	Maximum Wattage that PSU can supply to the system.

Table 48. IOM Status

Item	Description
Name/Location	IOM A
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Status	Off or On — Indicates whether the IOM is functioning.
Model	Model of the IOM.
Fabric Type	Networking type.
IP address	Only shows if IOM is turned on. This value is zero for a pass through type IOM.
Service Tag	The factory-assigned service tag.

Table 49. KVM Mapping Status

Item	Description
Server <number></number>	Displays a list of servers to which the KVM can be mapped.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Mapped	Displays a list of servers mapped to a KVM, if any.
Slot <number></number>	Indicates the server slot to which the KVM is mapped to. Possible values are SLOT-<01 to 04>.
Unmapped	Displayed if the KVM is not mapped to any of the servers.

Table 50. DVD Mapping Status

Item	Description
Server <number></number>	Displays a list of servers to which the DVD can be mapped.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Mapped	Displays a list of servers mapped to a DVD, if any.
Slot <number></number>	Indicates the server slot to which the DVD is mapped to. Possible values are SLOT-<01 to 04>.
Unmapped	Displayed if the KVM is not mapped to any of the servers.

Table 51. Blower Status

Item	Description
Name/Location	Example: Blower1, Blower2, and so on.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
RPM	Current blower speed in RPM.

Table 52. SPERC Status

Item	Description
SPERC: <number></number>	Displays the SPERC name in the format SPERC n, where 'n' is the SPERC number. Example: SPERC 1, SPERC 2, and so on.
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Working Status	On or Off — Indicates whether the SPERC is functioning.
Name: <name></name>	Name of the shared PERC. Example: SPERC
Health Status	Ok
Firmware Version	SPERC version
Manufacturer	Manufacturer name
State	Offline, Online, or Standby — Indicates the power status of a SPERC.

Table 53. PCIe Card Status

Item	Description
PCle Card <number></number>	Displays the PCle Card name in the format PCle Card <n>, where 'n' is the PCle Card number. Example: PCle Card 1, PCle Card 2, and so on.</n>
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Working Status	On or Off — Indicates whether the PCIe Card is functioning.
Name: <name></name>	Name of the PCIe Card.
Mapped to Server	Mapped or Unmapped.

Table 54. Hard Disk Drive Status

Item	Description
Hard Disk Drive: <number></number>	Displays the Hard disk drive name in the format Hard Disk Drive <n>, where 'n' is the hard drive number. Example: Hard Disk Drive 1, Hard Disk Drive 2, and so on.</n>
Error Messages	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related.
Power Status	Spun-Up, Transition, or Spun-Down — Indicates the power status of a hard disk drive
Manufacturer	Manufacturer name
Capacity	Available Storage capacity of the Hard Disk Drive in gigabytes (GB)
Firmware version	Firmware version of the Hard Disk Drive
State	Offline, Online, or Standby — Indicates the power status of the hard disk drive.

Table 55. Server Status

	Table 55. Server Status		
Item	Description		
Name/Location	Example: Server 1, Server 2, and so on.		
No Errors	If there are no errors, No Errors is displayed. Else, error messages are listed where critical ones are first listed, and then the warning-related. For more information, see "LCD Error Messages".		
Slot Name	Chassis slot name. For example, SLOT-01. i NOTE: You can set this table through the CMC CLI or CMC Web interface.		
Name	Name of the server, which the user can set through Dell OpenManage. The name is displayed only if iDRAC has finished booting, and the server supports this feature, else iDRAC booting messages are displayed.		
Model Number	Displays if iDRAC finished booting.		
Service Tag	Displays if iDRAC finished booting.		
BIOS Version	Server BIOS firmware version.		
Last POST Code	Displays the last server BIOS POST code messages string.		
iDRAC Firmware Version	Displays if iDRAC finished booting. i NOTE: iDRAC version 1.01 is displayed as 1.1. There is no iDRAC version 1.10.		
IP4 <enabled, disabled></enabled, 	Displays the current IPv4 enabled state.		
IP4 Address: <address, acquiring></address, 	Only displays if IPv4 is enabled.		
IP6 <enabled, disabled></enabled, 	Only displays if iDRAC supports IPv6. Displays current IPv6-enabled state.		
IP6 Local Address: <address></address>	Only displays if iDRAC supports IPv6 and IPv6 is enabled.		
IP6 Global Address: <address></address>	Only displays if iDRAC supports IPv6 and IPv6 is enabled.		
FlexAddress enabled on Fabrics	Only displays if the feature is installed. Lists the fabrics enabled for this server (that is, A, B, C).		

The information in the table is dynamically updated. If the server does not support this feature, then the following information does not appear, else Server Administrator options are as follows:

- Option "None" = No strings must be displayed on the LCD.
- Option "Default" = No Effect.
- Option "Custom" = Allows you to enter a string name for the server.

The information is displayed only if iDRAC has completed booting. For more information on this feature, see the RACADM Command Line Reference Guide for CMC in PowerEdge VRTX.

Frequently Asked Questions

This section lists the frequently asked questions about the following:

- RACADM
- Managing and Recovering a Remote System
- Active Directory
- FlexAddress and FlexAddressPlus
- IOM

Topics:

- RACADM
- Managing and Recovering a Remote System
- Active Directory
- FlexAddress and FlexAddressPlus
- IOM

RACADM

After performing a CMC reset (using the RACADM racreset subcommand), when a command is entered, the following message is displayed:

```
racadm <subcommand> Transport: ERROR: (RC=-1)
```

What does this message mean?

Another command must be issued only after CMC completes the reset.

Using the RACADM subcommands sometimes displays one or more of the following errors:

Local error messages — Problems such as syntax, typographical errors, and incorrect names. For example, ERROR:
 <message>

Use the RACADM help subcommand to display correct syntax and usage information. For example, if you have an error in clearing a chassis log, run the following sub-command.

```
racadm chassislog help clear
```

 CMC-related error messages — Problems where the CMC is unable to perform an action. The following error message is displayed:

```
racadm command failed.
```

To view information about a chassis, type the following command:

```
racadm gettracelog
```

While using firmware RACADM, the prompt changes to a ">" and the "\$" prompt is not displayed again.

If a non-matched double quotation mark (") or a non-matched single quotation (') is used in the command, the CLI changes to the ">" prompt and queues all commands.

To return to the \$ prompt, type <Ctrl>-d:

An error message Not Found is displayed while using the \$ logout and \$ quit commands.

Managing and Recovering a Remote System

Why are the remote RACADM and Web-based services unavailable after a property change?

It may take a minute for the remote RACADM services and the web interface to become available after the CMC Web server resets.

The CMC web server is reset after the following occurrences:

- Changing the network configuration or network security properties using the CMC web user interface.
- The cfgRacTuneHttpsPort property is changed (including when a config -f <config file> changes it).
- racresetcfg is used or a chassis configuration backup is restored.
- CMC is reset.
- A new SSL server certificate is uploaded.

My DNS server doesn't register my CMC?

Some DNS servers only register names with a maximum of 31 characters.

When accessing the CMC Web interface, a security warning stating that the SSL certificate was issued by a certificate authority that is not trusted is displayed.

CMC includes a default CMC server certificate to ensure network security for the web interface and remote RACADM features. This certificate is not issued by a trusted certificate authority. To address this security concern, upload a CMC server certificate issued by a trusted certificate authority (such as Thawte or Verisign).

Why is the following message displayed for unknown reasons?

Remote Access: SNMP Authentication Failure

As part of discovery, IT Assistant attempts to verify the device's **get** and **set** community names. In IT Assistant, the **get community name = public** and the **set community name = private**. By default, the community name for the CMC agent is public. When IT Assistant sends out a set request, the CMC agent generates the SNMP authentication error because it only accepts requests from **community = public**.

Change the CMC community name using RACADM. To see the CMC community name, use the following command:

racadm getconfig -g cfgOobSnmp

To set the CMC community name, use the following command:

racadm config -g cfgOobSnmp -o cfgOobSnmpAgentCommunity <community name>

To prevent SNMP authentication traps from being generated, enter input community names that are accepted by the agent. Since CMC only allows one community name, enter the same get and set community name for IT Assistant discovery setup.

When accessing the CMC Web interface, a security warning stating that the host name of the SSL certificate does not match the host name of CMC is displayed.

CMC includes a default CMC server certificate to ensure network security for the web interface and remote RACADM features. When this certificate is used, the web browser displays a security warning if the default certificate does not match the host name of CMC (for example, the IP address).

To address this security concern, upload a CMC server certificate issued to the IP address of CMC. When generating the certificate signing request (CSR) to be used for issuing the certificate, ensure that the common name (CN) of the CSR matches the IP address of CMC (for example, 192.168.0.120) or the registered DNS CMC name.

To ensure that the CSR matches the registered DNS CMC name:

- 1. In the left pane, click Chassis Overview.
- 2. Click Network.
 - The Network Configuration page appears.
- 3. Select the Register CMC on DNS option.
- 4. Type a CMC name in the DNS CMC Name field.
- 5. Click Apply Changes.

Active Directory

Does Active Directory support CMC login across multiple trees?

Yes. The CMC's Active Directory querying algorithm supports multiple trees in a single forest.

Does the login to CMC using Active Directory work in mixed mode (that is, the domain controllers in the forest run different operating systems, such as Microsoft Windows 2000 or Windows Server 2003)?

Yes. In mixed mode, all objects used by the CMC querying process (among user, RAC Device Object, and Association Object) must be in the same domain.

The Dell-extended Active Directory Users and Computers Snap-In checks the mode and limits users in order to create objects across domains, if in a mixed mode.

Does using CMC with Active Directory support multiple domain environments?

Yes. The domain forest function level must be in Native mode or Windows 2003 mode. In addition, the groups among Association Object, RAC user objects, and RAC Device Objects (including Association Object) must be universal groups.

Can these Dell-extended objects (Dell Association Object, Dell RAC Device, and Dell Privilege Object) be in different domains?

The Association Object and the Privilege Object must be in the same domain. The Dell-extended Active Directory Users and Computers Snap-In allows to create these two objects in the same domain only. Other objects can be in different domains.

Are there any restrictions on Domain Controller SSL configuration?

Yes. All SSL certificates for Active Directory servers in the forest must be signed by the same root certificate authority-signed certificate, because CMC only allows upload of one trusted certificate authority-signed SSL certificate.

The Web interface does not launch after a new RAC certificate is created and uploaded.

If Microsoft Certificate Services is used to generate the RAC certificate, the User Certificate option may have been used instead of Web Certificate, when creating the certificate.

To recover, generate a CSR, create a new Web certificate from Microsoft Certificate Services, and then upload it by running the following RACADM commands:

```
racadm sslcsrgen [-g] [-f {filename}]
racadm sslcertupload -t 1 -f {web sslcert}
```

FlexAddress and FlexAddressPlus

What happens if a feature card is removed?

There is no visible change if a feature card is removed. Feature cards can be removed and stored, or can be left in place.

What happens if a feature card that was used in one chassis is removed and put into another chassis?

The Web interface displays the following error message:

This feature card was activated with a different chassis. It must be removed before accessing the FlexAddress feature.

```
Current Chassis Service Tag = XXXXXXXXX

Feature Card Chassis Service Tag = YYYYYYYY

An entry is added to the CMC log that states:

cmc <date timestamp> : feature 'FlexAddress@YYYYYYYY' not activated; chassis ID='XXXXXXXXX'
```

What happens if the feature card is removed and a non-FlexAddress card is installed?

No activation or modifications to the card should occur. The card is ignored by CMC. In this situation, the **\$racadm featurecard -s** command returns the following message:

```
No feature card inserted ERROR: can't open file
```

If the chassis service tag is reprogrammed, what happens if there is a feature card bound to that chassis?

- If the original feature card is present in the active CMC on that or any other chassis, the Web interface displays the following error message:
 - This feature card was activated with a different chassis. It must be removed before accessing the FlexAddress feature.
 - O Current Chassis Service Tag = XXXXXXXX
 - Feature Card Chassis Service Tag = YYYYYYYY

The original feature card is no longer eligible for deactivation on that or any other chassis, unless Dell Service reprograms the original chassis service tag back into a chassis, and CMC that has the original feature card is made active on that chassis.

• The FlexAddress feature remains activated on the originally-bound chassis. The binding of that chassis feature is updated to reflect the new service tag.

Is an error message displayed if two feature cards are installed in a redundant CMC system?

A feature card in the active CMC is active and installed in the chassis. The second card is ignored by CMC.

Does the SD card have a write-protection lock on it?

Yes it does. Before installing the SD card into the CMC module, verify the write protection latch is in the unlock position. The FlexAddress feature cannot be activated if the SD card is write protected. In this situation, the **\$racadm feature -s** command returns this message:

No features active on the chassis. ERROR: read only file system

What happens if there is no SD card in the active CMC module?

The **\$racadm featurecard -s** command returns this message:

No feature card inserted.

What happens to FlexAddress feature if the server BIOS is updated from version 1.xx to version 2.xx?

The server module must be turned off before it can be used with FlexAddress. After the server BIOS update is complete, the server module does not get chassis-assigned addresses until the server is power cycled.

How can an SD card be recovered if the SD card was not in the chassis when the deactivation command was run on the FlexAddress?

The issue is that the SD card cannot be used to install FlexAddress on another chassis, if it was not in CMC when FlexAddress was deactivated. To recover use of the card, insert the card back into a CMC in the chassis that it is bound to, reinstall FlexAddress, and then deactivate the FlexAddress.

The SD card is properly installed and all the firmware or software updates are installed. The FlexAddress is active, the server deployment screen does not display the options to deploy it? What is wrong?

This is a browser caching issue. Log off from the browser and relaunch.

What happens to FlexAddress if I need to reset my chassis configuration using the RACADM command, racresetcfg?

The FlexAddress feature will still be activated and ready to use. All fabrics and slots are selected as default.

(i) NOTE: It is highly recommended that you turn off your chassis before running the RACADM command racresetcfg.

After disabling only the FlexAddressPlus feature (leaving FlexAddress still activated), why does the racadm setflexaddr command on the (stillactive) CMC fail?

If the CMC subsequently becomes active, with the FlexAddressPlus feature card still in the card slot, the FlexAddressPlus feature gets reactivated, and slot or fabric flexaddress configuration changes can resume.

IOM

After a configuration change, sometimes CMC displays the IP address as 0.0.0.0.

Click the **Refresh** icon to see if the IP address is set correctly on the switch. If an error is made in setting the IP/mask/gateway, the switch does not set the IP address and returns a 0.0.0.0 in all fields.

Common errors are:

- Setting the out-of-band IP address to be the same as, or on the same network as, the in-band management IP address.
- Entering an invalid subnet mask.
- Setting the default gateway to an address that is not on a network, which is directly connected to the switch.