

#####

DELL(TM) CHASSIS MANAGEMENT CONTROLLER (CMC)

#####

This document contains updated information about the Dell Chassis Management Controller (CMC).

For more information about CMC, including installation and configuration information, see the "Dell Chassis Management Controller Firmware Version 1.20 User's Guide" and the "Dell OpenManage(TM) Server Administrator User's Guide." These documents are located on the Dell Support website at "support.dell.com" and with your Product Documentation CD.

#####

TABLE OF CONTENTS

#####

This file contains the following sections:

- * Criticality
- * Minimum Requirements
- * Release Highlights
- * Known Issues for CMC v1.20
- * Known Issues for Documentation

#####

CRITICALITY

#####

2 - Recommended

#####

MINIMUM REQUIREMENTS

#####

The following subsections list operating systems that are compatible with the CMC.

=====

SUPPORTED SYSTEMS

=====

CMC is supported on the following Dell PowerEdge(TM) systems in the Dell PowerEdge M1000-e system enclosure:

- * Dell PowerEdge M600
- * Dell PowerEdge M605
- * Dell PowerEdge M805
- * Dell PowerEdge M905

=====

SUPPORTED WEB BROWSERS

=====

- * Microsoft(R) Internet Explorer 6.0 (32-bit) with SP1 for Windows 2000 Server family.
- * Microsoft Internet Explorer 6.0 (32-bit) with SP2 for Windows XP and Windows Server(R) 2003 family.
- * Microsoft Internet Explorer 7.0 for Windows Vista(R), Windows XP, and Windows Server 2003 family.
- * Mozilla Firefox 1.5 (32-bit).
- * Mozilla Firefox 2.0 (32-bit).

=====

FIRMWARE VERSIONS

=====

- * CMC Firmware Version: 1.20

=====

MODULE FIRMWARE VERSIONS RECOMMENDED

=====

Additional Chassis module Firmware recommended if CMC 1.20 is installed.

- * iDRAC Firmware Version: 1.20 (or later)
- * Server Module BIOS Firmware Version - M600: 2.04 (or later)
- * Server Module BIOS Firmware Version - M605: 2.06 (or later)
- * Server Module BIOS Firmware Version - M805: 1.0 (or later)
- * Server Module BIOS Firmware Version - M905: 1.0 (or later)

#####

RELEASE HIGHLIGHTS

#####

Fixes and Enhancements in 1.20

- * Support for PowerEdge 805 & 905 servers
- * CMC GUI Chassis Graphics
- * Enhanced Chassis Power reporting through the CMC
- * Enhanced CMC Firmware update to allow both CMC to be updated with single command.

FLEXADDRESS (Provided in CMC 1.10 Release)
#####

Required Module Firmware to use Chassis FlexAddress feature:

Component	Minimum required version
Ethernet Mezzanine card - Broadcom M5708t/M5708is,M5708	Firmware 4.4.1, iSCSI boot firmware 2.7.11, PXE firmware 4.4.3
FC Mezzanine card - QLogic QME2472	BIOS 2.04 or later
FC Mezzanine card - Emulex LPe1105-M4	BIOS 3.03a3 and firmware 2.72A2 or later
Blade Server BIOS	(PE M600) BIOS 2.02 or later (PE M605) BIOS 2.03 or later (PE M805) All BIOS versions (PE M905) All BIOS versions
iDRAC	Firmware 1.11 or later
CMC	Firmware 1.10 or later

- * FlexAddress: Prior to inserting the SD card into the CMC, the user must verify the write protection latch is in the "unlock" position. The FlexAddress feature cannot be activated if the SD card is write protected.
- * FlexAddress: The system BIOS must be upgraded prior to installing FlexAddress. If not, a warning icon will be displayed on the server health page. Once the system BIOS is updated, the server blade must be power cycled before the FlexAddress chassis assigned MAC addresses will be accepted by the server blade. The CMC will display chassis assigned MACs are configured but the server will be using the server assigned MAC configuration.
- * FlexAddress: If you issue a CMCCHANGEVER or RACRESET and then log into the CMC Web GUI, the FlexAddress webpage could take up to a min to update the switch configurations.
- * FlexAddress: If a chassis with a single CMC is downgraded with firmware prior to 1.10, the FlexAddress feature and configuration will be removed. Once the CMC firmware is upgraded to 1.10 or later, the FlexAddress feature will need to be reactivated and configured by the user.
- * FlexAddress: In a chassis with a two CMC, if replacing a CMC unit with one that has firmware prior to 1.10, the following procedure must be used to ensure the current FlexAddress feature and configuration will NOT be removed.
 1. Ensure the active CMC firmware is always version 1.10 or later
 2. Remove the standby CMC and insert the new CMC in its place.
 3. From the Active CMC, upgrade the standby CMC firmware to 1.10 or later.

Note: If you do not update the standby CMC firmware to 1.10 or later and a failover occurs FlexAddress feature will not be configured and the user will need to reactive the feature.

- * FlexAddress: Wake-On-LAN (WOL) requires BIOS to initialize MAC values. When the FlexAddress feature is deployed for the first time on a given blade, it requires at least one power-up and down sequence for FlexAddress to take effect. The reason for this is the FlexAddress on Ethernet devices is programmed by the BIOS. In order for the BIOS to program the address it would need to be functioning, this in turn requires blade to be powered up. Once the power-up and power-down sequence has been completed, FlexAddress would be available for Wake-On-LAN (WOL) function. Users may perform power-up and power-down sequence on the blade for fully deploying FlexAddress via iDRAC or CMC interface.
- * When changing from a Server-Assigned MAC to Chassis-Assigned MAC on Linux Based operating systems, additional configuration steps may be required.
 - o SLES 9 and SLES 10: Users may need to run YAST (Yet another Setup Tool) on their Linux system to configure their Network devices and then restart the network services.
 - o RHEL 4 and RHEL 5: Users will need to run Kudzu, a utility to detect and configure new/changed hardware on the system. Kudzu will present the user with The Hardware Discovery Menu, it will detect the MAC address change as Hardware was removed and new Hardware added.
- * Prior to installing FlexAddress, the user can determine the range of MAC addresses contained on their Flexaddress feature card by inserting the SD card into an USB "Memory Card Reader" and viewing the file "pwwn_mac.xml". The clear text XML file on the SD card will contain an XML tag "mac_start" which is the first starting hex MAC address that will be used for this unique MAC address range. The "mac_count" tag is the total number of MAC addresses that this SD card allocates. The total MAC range allocated can be determined by: "mac_start" + 0xCF (208 - 1) = mac_end.
Example:(starting_mac)00188BFFDCFA + 0xCF =(ending_mac)00188BFFDDC9

NOTE: You should lock the SD card prior to inserting in the USB "Memory Card Reader" to prevent accidently modifying any of the contents. You MUST UNLOCK the SD card before inserting into the CMC.

KNOWN ISSUES FOR BROWSERS
#####

- * In Internet Explorer Version 6, the log data may not display, leaving only the message, "Loading Chassis Event Log..." shown . To address this, go to Advanced Settings/Security and make sure the option, "Allow active content to run in files on My Computer" is NOT checked.
- * In Internet Explorer Version 6, if the security setting is set to restricted, the CMC User Interface on the Alert Management pages for

Email Alerts and SNMP Traps will pop up a Security Information message stating that the page contains both secure and non secure items and will ask if you want to continue. Please select "Yes". This is because the Internet Explorer Version 6 does not allow the use of hidden IFRAMES on secure (SSL) pages. (183022)

* In Firefox Version 1.5, you must manually refresh pages. Automatic page updating is not fully supported in this version of the Firefox browser.

* In Internet Explorer Version 6, after updating the Active CMC you may need to close the browser used to login to the CMC before attempting to login again. (232942)

KNOWN ISSUES FOR CMC
#####

* CMC firmware 1.2 has enhanced the power allocation algorithms to allow blade servers to receive higher power allocations. If a user downgrades the CMC to 1.10 or earlier firmware version, blades with the new higher power allocations could be powered off because the earlier firmware cannot support the higher chassis power allocations. If this occurs, you must power the blade back on. (230143)

* After a CMC reset, the CMC may require up to 30 seconds after the serial login prompt is displayed before RACADM commands will be accepted. Commands issues prior to that time may receive the message "ERROR: Unable to perform requested operation."

* RACADM CLI based command line tool uses TFTP or FTP to transfer image files for all firmware updates. For TFTP transfers only, the default port (69) is supported for transfers with version 1.20. (157754)

* Clearing the CMC Log can take a long time. Please allow up to one (1) minute for this operation to complete.(152860)

* If the CMC is on a private network without access to the Internet and you are using Internet Explorer 6 SP 2 or Internet Explorer 7, you may experience delays of up to 30 seconds when using remote RACADM commands. (161019)

* Some USB-to-serial adapters have been found to generate a large number of spurious interrupts when plugged in. If the adapter is connected to the CMC's serial port when this happens, the CMC can become overloaded when attempting to service these interrupts and may reboot. This problem is exacerbated when the serial cable is very long, causing voltage levels to drop and noise on the serial line to increase. To avoid this issue, Dell recommends first connecting the USB-to-serial adapter into the USB port, before connecting to the CMC. Dell also recommends disconnecting the adapter from the CMC before rebooting or performing other power management functions on a system that is attached to the CMC. (180373)

* The "connect" to server command enables the user to access the server's serial port through CMC's serial port. After this connection the user will be able to see the server's console redirection through CMC's serial port.

Following are the prerequisites for this feature:

1. The server should be Powered on.
2. Serial redirection has to be enabled in the servers BIOS Setup.
3. Not all keys will work so the user has to provide appropriate escape sequences for "CTRL+ALT+DEL" etc. The initial redirection screen displays the necessary escape sequences.(184385)

* If you setup Active Directory (AD) on the CMC using extended schema and the built-in Administrator privilege object and then attempt to login to the CMC User Interface using this AD account, after successful AD login, the user name and privilege level displayed on the right hand side of the User Interface just beneath log out link is displayed as a custom user rather than the privilege as created on the AD side (example: Administrator, power user). (183449)

* Using the RACADM command line utility if you attempt to set the DNS CMC Name or DNS Domain Name without the proper rules (Rules: start with an alphabetic character (a-z, A-Z) and follow by an alphanumeric (a-z, A-Z, 0-9) or a valid symbol (such as -)) then the utility will display a non-specific error message (ERROR: Unable to perform requested operation). Please enter a valid name for the above mentioned names. (173204)

* If the standby CMC firmware was updated at a previous time and later made active, you will see an entry in the CMC log (raclog) that states the firmware update was successful but with the most recent time stamp. This is because the standby CMC cannot register its entries to the logs until it becomes active. (186858)

* If using 'racadm fwupdate' to update a second component firmware image, the user should check the status of the first update using the 'racadm fwupdate -s' command prior to attempting a second firmware update. (230487)

* A 'racadm config' operation may fail, due to configuration property definition changes across Firmware versions. For example, if the set of allowable values for a configuration property has been changed, and a snapshot of the prior values (from 'racadm getconfig') is used in a 'racadm config' operation on a newer version of firmware, the prior values may no longer be accepted, and thus cause the racadm config operation to fail.(229764)

To resolve this issue, comment out the prior value in the captured file, and restart the 'racadm config' operation.

The following lists the racadm property definitions that differ, depending on the CMC firmware version:

group: cfgNetTuning, object: cfgNetTuningNicSpeed
- CMC Firmware 1.0 and 1.10:
Allowed values: 10, 100, 1000 (default 1000)

- CMC version 1.20 and later:
Allowed values: 10, 100 (default 100)

* Modification of configuration objects cfgChassisRedundancyPolicy and cfgChassisDynamicPSUEngagementEnable in the cfgChassisPower group can in rare instances appear successful via the GUI and RACADM CLI, but actually fail to take effect. To determine if a modification has failed, review the raclog messages for the following entry:

Redundancy policy change denied - Power Management busy, retry after xy seconds.

To correct this issue, re-enter the modifications.

* During a server module power cycle, iDRAC may log a link tuning error that will cause an amber LED and a red "X" on the iDRAC health page. To determine if a link tuning error has occurred, review the iDRAC's SEL log for the following entry:

LinkT/FlexAddr: Link tuning sensor failed to get, link tuning or flex address data from BMC/iDRAC

To correct this issue, perform a server module power cycle. (230455)

KNOWN ISSUES FOR USER INTERFACE ONLINE HELP
#####

This section provides additional information about known issues with the CMC Firmware version 1.20 User Interfaces online help.

* The help for the Chassis - Power Management - Budget Status - System Power Policy configuration section in the localized languages should read as follows:

System Input Power Cap: Displays the user defined maximum power consumption limit for the entire system (chassis, CMC, servers, I/O modules, power supply units, iKVM, and fans). The CMC will enforce this limit via throttling or by powering off lower priority blades. The value for system input power cap is displayed in watts, BTU/h and percent units.

If the chassis power consumption exceeds the System Input Power Cap, then the performance of lower priority servers is reduced until total power consumption falls below the cap.

In cases where the servers are set to the same priority, then the selection of the server for throttling, or power-off action, is based on the server slot number order. For example, the server in slot 1 is selected first and the sever in slot 16 selected last.

Surplus for Peak Performance: The surplus for peak performance value is the difference between the System Input Power Cap and the sum of the Maximum Input Power Allocated to Servers and the Maximum Input Power Allocated to

Chassis Infrastructure. The value for surplus for peak performance is indicated in both the watts and BTU/h units.

Redundancy Policy: Displays the current redundancy configuration: AC Redundancy, Power Supply Redundancy, or No Redundancy.

* AC Redundancy: Power input is load-balanced across all supplies. Three of the power supplies are connected to one AC grid, while the other three are connected to another AC grid. When the system is running optimally in AC Redundancy mode, power is load-balanced across all active supplies. In case of failure, the power supplies on the functioning AC grid take over at 100% capacity.

NOTE: In AC Redundancy mode, a difference in the number of power supplies between the two AC circuits (for example, three power supplies on one AC circuit and two on the other AC circuit) will cause degradation in the redundancy.

* Power Supply Redundancy: The capacity of the highest-rated power supply in the chassis is kept as a spare, ensuring that a failure of any one power supply will not cause the server modules or chassis to power-down.

Power Supply Redundancy mode does not use all six power supplies; it uses a maximum of four power supplies. Power supplies in excess of four do not participate in Power Supply Redundancy unless there is a failure or removal of one of the power supplies.

* No Redundancy: Power from all three power supplies on one AC circuit (grid) is used to power-on the entire chassis, including the chassis, servers, I/O modules, iKVM, and CMC.

NOTE: The No Redundancy mode uses only three power supplies at a time, without backup. Failure of one of the three power supplies being used could cause the server modules to lose power and data.

Dynamic Power Supply Engagement: Displays whether Dynamic Power Supply Engagement is enabled or disabled. Enabling this feature allows the CMC to put under-utilized power supplies into standby mode based on the redundancy policy that is set and the power requirements of the system. Putting under-utilized power supplies into standby mode increases the utilization of active power supplies and thus the efficiency of the online supplies, saving power. (234463)

* The help for the Chassis - Power Management - Configuration - Configuring Power Budget and Redundancy section in the localized languages should read as follows:

* Surplus for Peak Performance: The surplus for peak performance value is the difference between the System Input Power Cap and the Power Required for Peak Performance (sum of the Maximum Input Power Allocated to Servers and the Maximum Input Power Allocated to Chassis Infrastructure). The value for surplus for peak performance is indicated in both the watts and BTU/h units.

* Power Required for Peak Performance: The power required for peak performance value is the sum of the Maximum Input Power Need of Servers that are powered on and the Input Power Allocated to Chassis Infrastructure. The value for power required for peak performance is indicated in both the watts and BTU/h units. If System Input Power Cap is set to less than Power Required for Peak Performance, some servers may throttle under extreme load.

* **System Input Power Cap:** System Input Power Cap is the max AC power that the system is allowed to allocate to servers and chassis infrastructure. It can be configured by the user to any value that exceeds the minimum power needed for Servers that are powered on and the chassis Infrastructure. If an attempt is made to configure a value that falls below the minimum power needed for Servers and Chassis Infrastructure, the attempt would fail. The power allocated to Servers and Chassis Infrastructure can be found in the User Interface on the Chassis -> Power Management-> Power Budget status page under Power Budgeting section or via CLI RACADM utility command (racadm getpbinfo). Users can power OFF one or more server(s) to lower the current Power allocation, and re-attempt setting a lower value for System Input Power Cap (if desired) or simply configure the cap prior to powering on the servers.

To change this setting, it is possible to enter a value in any of the units. The interface ensures that the unit field that was last changed will be the value that is submitted when those changes are applied.

NOTE: Refer to the Datacenter Capacity Planner (DCCP) tool at www.dell.com/calc for capacity planning.

NOTE: When value changes are specified in watts, the submitted value will exactly reflect what is actually applied. However, when the changes are submitted in either of the BTU/h or percent units, the submitted value may not exactly reflect what is actually applied. This is because these units are converted to watts and then applied and the conversion will be susceptible to some rounding error.

* **Redundancy Policy** This option will allow you to select one the following options:

* **No Redundancy:** Power from all three power supplies on one AC circuit (grid) is used to power-on the entire chassis, including the chassis, servers, I/O modules, iKVM, and CMC.

NOTE: The No Redundancy mode uses only three power supplies at a time. If 3 PSUs are installed, then there is no backup available. Failure of one of the three power supplies being used could cause the servers to lose power and/or data. If PSUs 4-6 are present, then these will become redundant and will become available in the event of an online PSU going down.

* **Power Supply Redundancy:** The capacity of the highest-rated power supply in the chassis is kept as a spare, ensuring that a failure of any one power supply will not cause the server modules or chassis to power-down (hot spare).

Power Supply Redundancy mode does not utilize all six power supplies, but rather a maximum of four and a minimum of two power supplies. Power Supply Redundancy mode prevents server modules from powering up if the power consumption of the chassis exceeds the rated power. Failure of two power supplies may cause some or all server modules in the chassis to power-down. Server modules are not throttled in this mode.

* **AC Redundancy:** This divides the 6 PSUs into two power grids (PSUs 1-3 making up power grid 1 and PSUs 4-6 making up power grid 2). 6 PSUs are required to have a fully redundant AC Redundancy power policy. In this configuration 3 PSUs in one grid will be online and 3 PSUs in the other grid will be redundant. Failover will occur when any of the 3 PSUs in the online grid fail. This will cause the redundant PSUs to convert to online and will report the redundancy policy as degraded.

NOTE: In AC Redundancy mode, a difference in the number of power supplies between the two AC circuits (for example, three power supplies on one AC circuit and two on the other AC circuit) will cause a degradation in the redundancy.

- * Enable Dynamic Power Supply Engagement: Enables (when checked) dynamic power management. In Dynamic Engagement mode, the power supplies are turned ON or OFF based on power consumption, optimizing the energy consumption of the entire chassis.

For example, your power budget is 3000 watts, your redundancy policy is set to AC redundancy mode, and you have four power supply units. The CMC determines that two of the power supply units can manage the AC redundancy while the other two remain in standby mode. If an additional 2000W of power is needed for newly installed servers, then the two standby power supply units are engaged. The two standby power supply units are also engaged in the event of a power failure.

- * Disable Chassis Power Button: Disables (when checked) the chassis power button. If the checkbox is checked and the user attempts to change the power state of the chassis through the chassis power button, the user action is ignored. (234463)

KNOWN ISSUES FOR DOCUMENTATION
#####

No known issues for this release.

#####

Information in this document is subject to change without notice.
(C) 2008 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: "Dell", "Dell OpenManage", and "PowerEdge" are trademarks of Dell Inc.; "Microsoft", "Windows", "Windows Vista", "Windows Server", and "Active Directory" are trademarks or registered trademarks of Microsoft Corporation.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

August 2008