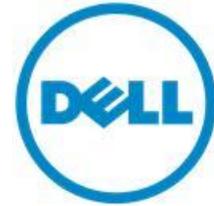


PowerEdge M1000e



Technical Guide



Built from the ground up to combat data center sprawl and IT complexity, the PowerEdge M1000e enclosure delivers one of the most energy-efficient, flexible and manageable blade server products on the market.

This document is for informational purposes only. Dell reserves the right to make changes without further notice to any products herein. The content provided is as is and without express or implied warranties of any kind.

Dell, the DELL logo, PowerEdge, PowerConnect, RapidRails, VersaRails, FlexAddress, EqualLogic, PowerVault, OpenManage, and Compellent are trademarks of Dell, Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows Server, and Active Directory are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

©Copyright 2013 Dell Inc. All rights reserved. Reproduction or translation of any part of this work beyond that permitted by U.S. copyright laws without the written permission of Dell Inc. is unlawful and strictly forbidden.

September 2013 | Version 3.0



Table of contents

1	Blade enclosure overview.....	5
	Introduction.....	5
	Key technologies	6
2	Blade enclosure specifications	8
3	Blade enclosure views and features	12
	Front view and features	12
	Back view and features.....	14
	Power supply indicators	18
	Security features	18
4	Midplane	20
5	Processors and Memory	23
6	Storage	24
7	Networking and Input/Output.....	25
	Embedded NICs/LAN on Motherboard	25
	Input/Output	25
8	Power, thermal and acoustics	35
	Power supplies	35
	Power supply specifications	38
	Power redundancy options	38
	Power management.....	46
	Heat dissipation.....	48
	Acoustics	50
9	Rack information	52
	Rails	53
	Cabling.....	55
	Rack and Stack.....	57
10	Virtualization	58
11	Systems management.....	59
	Server management.....	60
	Enclosure management	60
	Integrated keyboard and mouse controller.....	64
Appendix A.	Additional specifications.....	66
	Enclosure dimensions and weight	66
	Environmental specifications	66
	Video specifications	66
	Peripherals	66
Appendix B.	Standards compliance	68
Appendix C.	Additional resources.....	70

Tables

Table 1.	Technical specifications.....	9
Table 2.	Fabric specifications	28
Table 3.	FlexAddress features and benefits.....	33
Table 4.	2360W power supply.....	37
Table 5.	2700W power supply.....	38
Table 6.	Supported power supply redundancy options	39
Table 7.	Grid redundancy.....	40



Table 8.	Power supply redundancy	40
Table 9.	No Redundancy	40
Table 10.	Power supply compatibility per CMC revision	41
Table 11.	Quarter-height servers per power supply configuration	42
Table 12.	Half-height servers per power supply configuration	44
Table 13.	Full-height servers per power supply configuration	45
Table 14.	Supported racks	55
Table 15.	Typical modular server system rack height and cable reduction	55
Table 16.	Dimensions and weight	66
Table 17.	Industry standard documents	68
Table 18.	Additional resources	70

Figures

Figure 1.	Server density comparison	8
Figure 2.	Front of M1000e chassis	12
Figure 3.	M1000e LCD panel in recessed position	13
Figure 4.	M1000e LCD panel during usage	13
Figure 5.	LCD panel capabilities	14
Figure 6.	Back of enclosure	15
Figure 7.	Installed CMC, I/O module, and power supply blanks	16
Figure 8.	Simplified cabling	17
Figure 9.	Power supply indicators	18
Figure 10.	Midplane	20
Figure 11.	M1000e midplane front view	21
Figure 12.	M1000e midplane back view	22
Figure 13.	High-speed I/O architecture	27
Figure 14.	Ethernet growth path	29
Figure 15.	Difference between pass-through and switch modules	30
Figure 16.	FlexAddress addresses	32
Figure 17.	FlexAddress screen in the CMC	33
Figure 18.	Power supplies in M1000e	35
Figure 19.	M1000e power supply	36
Figure 20.	Power architecture	37
Figure 21.	Grid redundancy	39
Figure 22.	PMBus communication channels	47
Figure 23.	Server cooling air profile	48
Figure 24.	I/O module cooling air profile	49
Figure 25.	Power supply inlet and cooling air profile	50
Figure 26.	M1000e mounted in a rack	52
Figure 27.	Empty enclosure in a rack	53
Figure 28.	RapidRails static rails	54
Figure 29.	VersaRails static rails	54
Figure 30.	M1000e strain relief bar and cable enumerator clip (12 per kit)	56
Figure 31.	M1000e strain relief bar and cable enumerator clips	56
Figure 32.	Examples of major virtualization platforms supported	58
Figure 33.	System Management Architecture Diagram	59
Figure 34.	Web interface for the chassis management controller	61
Figure 35.	Chassis management controller	62
Figure 36.	M1000e iKVM	64
Figure 37.	Front keyboard/video ports	65



1 Blade enclosure overview

Introduction

The Dell™ PowerEdge™ M1000e Modular Blade Enclosure is the rock-solid foundation for Dell's blade server architecture, providing an extremely reliable, flexible and efficient platform for building any IT infrastructure. The M1000e enclosure is built from the ground up to combat data center sprawl and IT complexity, delivering one of the most energy efficient, flexible and manageable blade server implementations on the market.

Leading energy efficiency

The M1000e enclosure takes advantage of its world-class design by coupling ultra-efficient power supplies with large variable-speed fans and optimized airflow to effectively cool the entire chassis while using less power.

Effortless scalability

Only Dell provides complete, scale-on-demand switch designs. With additional I/O slots and switch options, you have the flexibility you need to meet ever-increasing demands for I/O consumption. Plus, Dell's FlexIO modular switch technology lets you easily scale to provide additional uplink and stacking functionality — no need to waste your current investment with a “rip and replace” upgrade.

Easy-to-use, powerful management tools

The M1000e helps reduce the cost and complexity of managing computing resources so you can focus on growing your business or managing your organization with features such as:

- Centralized Chassis Management Controller (CMC) modules for redundant, secure access paths for IT administrators to manage multiple enclosures and blades from a single interface.
- Dynamic and granular power management so you have the capability to set power thresholds to help ensure your blades operate within your specific power envelope.
- Real-time reporting for enclosure and blade power consumption, as well as the ability to prioritize blade slots for power, providing you with optimal control over power resources.

FlexAddress technology: the simple, low-cost way to limit downtime

Dell's patent-pending FlexAddress™ technology allows any M-Series blade enclosure to assign the World Wide Name (WWN) or Media Access Control (MAC) address of Fibre Channel, Ethernet and iSCSI controllers to an M1000e blade slot instead of directly to the blade. By removing the network and storage identity from the server hardware, you can upgrade and replace components or the entire blade server without having to change the identity on the network or rezoning switches. Unlike other solutions, which often require separate management interfaces and proprietary hardware, FlexAddress works with any network and is implemented directly from the integrated CMC by simply selecting the chassis slots and fabrics that you want to enable. FlexAddress delivers persistent network and storage identities, equipping your data center to handle predictable or even unplanned changes — add, upgrade, or remove servers without affecting your networks.



Key technologies

The PowerEdge M1000e is designed to help you be more efficient with time, power and cooling, investment, and system performance. It is a breakthrough Dell-engineered and patent-pending design that maximizes flexibility, power and thermal efficiency, system-wide availability, performance and manageability. The chassis integrates the latest in management, I/O, and power and cooling technologies in a modular, easy-to-use package. Designed from the ground up to support current and future generations of server, storage, networking and management technologies, the PowerEdge M1000e includes the headroom necessary to scale for the future.

Dell optimized the PowerEdge M1000e modular server enclosure and server modules to:

- Maximize flexibility — modular I/O, power, cooling and management architecture
- Maximize longevity — optimized power and cooling design supports current and future generations of server modules and I/O; I/O bandwidth to support not only today's generation of 10Gb Ethernet, 20Gbps InfiniBand® and 4Gbps Fibre Channel, but up to 40Gbps QDR InfiniBand, 10Gbps Serial Ethernet, and 8Gbps Fibre Channel
- Lower TCO — lower cost than rack-mount servers with equivalent features; best-in-class power and cooling efficiency

The PowerEdge M1000e modular server enclosure solution supports server modules, network, storage, and cluster interconnect modules (switches and pass-through modules), a high-performance, and highly available passive midplane that connects server modules to the infrastructure components, power supplies, fans, and integrated KVM and CMC. The PowerEdge M1000e uses redundant and hot-pluggable components throughout to provide maximum uptime.

The M1000e provides identical and symmetrical fabric options B and C for each modular server. Ethernet I/O switches support I/O sub-modules that provide external I/O flexibility of stacking ports, 10GbE copper ports or 10GbE optical ports. True modularity at the system and subsystem level provides simplicity of extension and enhancement, now and in the future.

The main benefits of these features include improved:

- Data center density
- Power and cooling efficiency
- Flexibility
- Scalability
- Virtualization capability
- Ease of deployment
- Manageability

Together, these factors help you do more with your server investment.

Virtually unlimited in scalability, the PowerEdge M1000e chassis provides ultimate flexibility in server processor and chipset architectures. The M1000e infrastructure can simultaneously support both Intel and AMD server architectures, while cutting-edge mechanical, electrical and software interface definitions enable multi-generational server support and expansion.

The chassis features:

- A high-speed passive midplane that connects the server modules in the front and to power, I/O and management infrastructure in the back of the enclosure



- Comprehensive I/O options to support dual links of 56 Gbps (with 4x FDR InfiniBand), which provide high-speed server module connectivity to the network and storage now and well into the future
- Thorough power-management capabilities, including delivering shared power to ensure full capacity of the power supplies available to all server modules
- Broad management ability, including private Ethernet, serial, USB and low-level management connectivity between the CMC, keyboard, video and mouse (KVM) switch, and server modules
- Up to two chassis management controllers (CMC-1: standard; CMC-2: providing optional redundancy) and one optional integrated KVM (iKVM) switch
- Up to six hot-pluggable, redundant power supplies and nine hot-pluggable, N+1 redundant fan modules
- System front control panel with LCD panel, two USB keyboard and mouse connections and one video crash-cart connection
- Support for Dell EqualLogic™ PS-M4110 Blade Array with seamless integration into the M1000e chassis



2 Blade enclosure specifications

The Dell PowerEdge M1000e modular server enclosure is a breakthrough in enterprise server architecture. The enclosure and its components spring from a revolutionary, ground-up design incorporating the latest advances in power, cooling, I/O and management technologies. These technologies are packed into a highly available rack-dense package that integrates into standard Dell and third-party 2000mm depth racks.

The M1000e supports up to 8 full-height, 16 half-height, or 32 quarter-height blade server modules. The chassis guide and retention features are designed so that alternative module form factors are possible. The chassis architecture is flexible enough for various types of front-loading modules such as servers and storage.

Dell server modules are accessible from the front of the M1000e enclosure. At the bottom of the enclosure is a flip-out, multiple-angle LCD screen for local systems management configuration, system information and status. The front of the enclosure also contains two USB connections for USB keyboard and mouse, a video connection, and a system power button. The USB and video ports on the front control panel work only when the iKVM module is installed — the iKVM provides the capability to switch the KVM between the blades.

For more information on the M1000e blade enclosure features, see the *Dell PowerEdge M1000e Enclosure Owner's Manual* on Dell.com/Support/Manuals.

Using 32 quarter-height blades in the PowerEdge M1000e provides up to 220% more computational density when compared to 1U servers.

Figure 1. Server density comparison



Greater density means:

- Smaller footprint
- More processing performance
- More RAM capacity
- Lower power consumption per unit
- Easier manageability



Table 1 lists the technical specifications for the PowerEdge M1000e blade enclosure. For the latest information on supported features, visit Dell.com/PowerEdge.

Table 1. Technical specifications

Feature	Parameter
Chassis enclosure	<p>Form factor: 10U modular enclosure holds up to 8 full-height, 16 half-height, or 32 quarter-height blade servers</p> <p>Dimensions: 44.0cm (17.3 in.) H x 44.7cm (17.6 in.) W x 75.4cm (29.7 in.) D</p> <p>Weight:</p> <ul style="list-style-type: none"> • Empty chassis – 98 lb • Chassis with all rear modules (I/O modules, PSUs, CMCs, KVM) – 176 lb • Chassis fully loaded with blade servers and rear modules – 394 lb
Power supplies	<p>Up to six 2700W high-efficiency hot-plug power supplies or up to six -48V DC 2700W power supplies</p> <p>Based on Dell's Energy Smart technologies, the M1000e power supplies deliver greater levels of efficiency, even at low levels of utilization.</p> <p>Supported power supply configurations include:</p> <ul style="list-style-type: none"> • 3+3 and 2+2 (grid redundancy) • 2+1, 3+1, 4+2, and 5+1 (power supply redundancy) • 2+0 and 3+0 (non-redundant mode) <p>The M1000e chassis supports dynamic power supply engagement (DPSE) mode, which if enabled, puts lightly loaded power supplies into standby mode, driving up the utilization and the efficiency of the active supplies.</p> <p>Dell supports either 110–120V or 208–240V AC power supply input and supports a wide range of power distribution options. Dell recommends 208–240V AC for all production environments.</p>
Cooling fans	<p>M1000e chassis comes standard with nine hot-pluggable, redundant fan modules:</p> <ul style="list-style-type: none"> • Based on Dell Energy Smart technologies, M1000e fans are a breakthrough in power and cooling efficiency • Fans deliver low-power consumption, but also use next-generation fan technologies to ensure the lowest possible amount of fresh air is consumed to cool the enclosure
Input device	<p>Front control panel with interactive graphical LCD</p> <ul style="list-style-type: none"> • Supports initial configuration wizard • Local server blade, enclosure, and module information and troubleshooting <p>Two USB keyboard/mouse connections and one video connection (requires the optional Avocent® iKVM switch to enable these ports) for local front "crash cart" console connections that can be switched between blades</p>



Feature	Parameter
Enclosure I/O modules	<p>Up to six total blade I/O modules for three fully redundant fabrics that take advantage of a fully passive midplane that can scale up to 8.4Tbps with current technologies featuring 1/10Gb Ethernet to servers, 40Gb Ethernet to top-of-rack switches, Fibre Channel 8Gb, and InfiniBand QDR/FDR10/FDR. Dell's blade Ethernet devices also have FlexIO technology, providing on-demand uplink scalability and investment protection unrivaled in the blade server market.</p> <p>Dell PowerEdge M I/O Aggregator Supports Active System Manager and CMC Aggregator GUI. 1/10 Gb Ethernet connectivity with Zero touch FCoE and converged iSCSI deployment. 32 internal 10GE ports provide full redundancy to the M420 blade on fabric A.</p> <p>Dell Networking MXL 10/40Gb Ethernet Switch Converged 1/10/40Gb Ethernet switch with up to 56-10GE ports (32 internal), converged iSCSI and FCoE (transit to ToR), 320Gb stacking of up to 6 switches, FlexIO technology, and PVST+ industry standard spanning tree. 32 internal 10GE ports provide full redundancy to the M420 blade on fabric A.</p> <p>Dell PowerConnect™ M8024-k 10Gb Ethernet Switch Converged 1/10Gb Ethernet switch with up to 24-10GE ports (16 internal), converged FCoE (transit to ToR), stacking of up to 6 modules, FlexIO technology, and simple mode.</p> <p>Dell M8428-k 10Gb Converged Network Switch Converged 1/10Gb Ethernet switch with 24-10GE Ethernet ports (16 internal) and 4 native Fibre Channel ports, converged iSCSI and FCoE (FC breakout at blade switch). Fibre Channel ports support 2Gb, 4Gb, and 8Gb.</p> <p>Dell PowerConnect M6220 Gigabit Ethernet Switch 1Gb Ethernet switch with up to 24 ports (16 internal), stacking of up to 12 switches, simple mode, and 4 x fixed copper 10/100/1000Mb Ethernet uplinks and FlexIO technology for up to four additional 10GbE uplinks.</p> <p>Dell PowerConnect M6348 Gigabit Ethernet Blade Switch 1Gb Ethernet switch with up to 48 ports (32 internal), stacking of up to 12 modules, simple mode, and 16 external fixed 10/100/1000Mb Ethernet RJ-45 ports to support 32 internal server GbE connections supplied by quad-port Gigabit Ethernet adapters plus up to four 10Gb uplink ports.</p> <p>Dell 10Gb Gigabit Ethernet Pass-Through -k Module Supports 16 x 10Gb connections only.</p> <p>Cisco® Nexus™ B22DELL Fabric Extender (FEX) 16 internal 10GbE ports and 8 external 10GbE ports; provides 10GbE connectivity to Cisco Nexus 5500 series infrastructure.</p> <p>Cisco Catalyst® Blade Switch M 3032 Includes 4 x fixed copper 10/100/1000Mb Ethernet uplinks standard plus two optional module bays, supporting 2 x 1Gb copper or optical SFPs each.</p> <p>Cisco Catalyst Blade Switch M 3130G Includes 4 x fixed copper 10/100/1000Mb Ethernet uplinks, 64Gb full duplex StackWise Plus stacking ports, and support for two x2 modules for up to a total of two 10Gb CX4 or SR/LRM uplinks.</p> <p>Cisco Catalyst Blade Switch M 3130X Includes 4 x fixed copper 10/100/1000Mb Ethernet uplinks, 64Gb full duplex StackWise Plus stacking ports, and support for two x2 modules for up to a total of two 10Gb CX4 or SR/LRM uplinks.</p> <p>Brocade® M5424 8Gb Fibre Channel Switch Includes eight external SAN ports.</p> <p>Dell 8Gb/4Gb Fibre Channel SAN Module NPIV Port Aggregator</p> <p>Emulex® 8Gb or 4 Gb/s Fibre Channel Pass-Through Module</p> <p>Mellanox® M4001Q QDR, M4001T FDR10, M4001Q FDR InfiniBand Switches 16 internal ports and 16 external QSFP+ ports in a single wide form factor.</p> <p>Brocade M6505 16Gb Fibre Channel SAN I/O Module 24 x 16Gbps Fibre Channel ports (16 internal server ports and 8 external fabric ports). Zero footprint, hot-pluggable design with no additional fans or power supplies. Dynamic Ports on Demand (PoD) and "pay-as-you-grow" port upgrades for 12-port configurations.</p>



Feature	Parameter
Systems management	<p data-bbox="402 201 1430 264">One (standard) or a second optional (redundant) Chassis Management Controller(s) (CMC) that provides:</p> <ul data-bbox="402 275 1463 1052" style="list-style-type: none"> • Single secure interface for inventory, configuration, monitoring and alerting for the chassis and all components • Multi-chassis management capability allows up to nine chassis and 288 servers to be managed from a single, embedded, agent-free interface • Automated and embedded one-to-many blade BIOS and firmware updates, independent of the operating system through iDRAC with Lifecycle Controller • One-to-many blade server BIOS capture and replication • Real-time power/thermal monitoring and management • Real-time system AC power consumption with resettable peak and minimum values • System-level power limiting and slot-based power prioritization • DPSE functionality that can help lower overall system power consumption by ensuring power supplies run at their optimal efficiency points • Fan speed control using Dell's enhanced efficiency technologies to ensure fans are delivering optimal cooling while minimizing power consumption and airflow • Secure web (SSL) and command line (Telnet/SSH) interfaces • Support for multiple levels of user roles and permissions, including integration into Microsoft® Active Directory® services and lightweight directory access protocol (LDAP) services for authentication • 2 x 10/100/1000Mb Ethernet ports + 1 serial port • Single point of connection from the management network to iDRAC with Lifecycle Controller on each of the blades and the management interfaces on the integrated I/O modules • 2nd Ethernet port that supports daisy chaining of CMCs for improved cable management <p data-bbox="402 1062 1300 1094">Optional Integrated Avocent® keyboard, video, and mouse (iKVM) switch</p> <ul data-bbox="402 1104 1268 1136" style="list-style-type: none"> • Enables USB keyboard/mouse and video port on front control panel <p data-bbox="402 1146 1409 1209">Dell OpenManage™ Systems Management tools (available as software downloads from dell.com)</p> <ul data-bbox="402 1220 1442 1377" style="list-style-type: none"> • Dell OpenManage Server Administrator — monitoring agents and 1:1 management utilities • Dell OpenManage Essentials — one-to-many remote discovery, inventory, monitoring, and update for M1000e chassis and blades • Integration and Connections with third-party management solutions
Additional storage options	<p data-bbox="402 1392 971 1423">Internal: Dell EqualLogic PS-M4110 Blade Array</p> <p data-bbox="402 1434 1317 1465">External: Dell Compellent™, Dell EqualLogic, and Dell PowerVault™ storage</p>
Rack support	<p data-bbox="402 1476 1442 1537">RapidRails™ static rails for 4-post square hole racks; VersaRails™ static rails for 4-post square or unthreaded round hole racks</p>



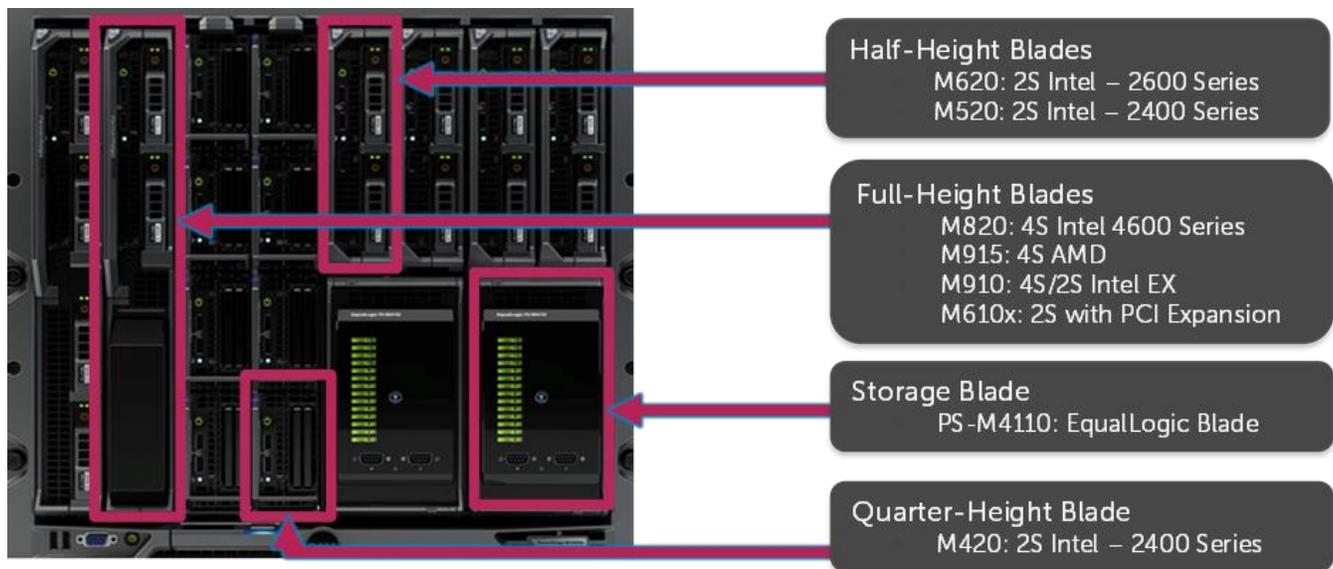
3 Blade enclosure views and features

The Dell M1000e supports up to 32 quarter-height, 16 half-height, or 8 full-height server modules. The chassis guide and retention features are designed such that alternative module form factors are possible. The chassis architecture is flexible enough that server, storage, or other types of front-loading modules are possible.

Front view and features

The server modules the M1000e enclosure supports occupy a slot accessible in the front of the enclosure. The enclosure has also been designed to accommodate other form factors, including half-height, double-wide modules such as the Dell EqualLogic PS-M4110 Blade Array. Server modules can be freely located within each quadrant dependent on the blade form factor.

Figure 2. Front of M1000e chassis



At the bottom of the enclosure is a flip out multiple angle LCD screen for local systems management configuration, system information, and status. The front of the enclosure also contains two USB connections for USB keyboard and mouse, a video connection and the system power button. The front control panel's USB and video ports work only when the iKVM module is installed, as the iKVM provides the capability to switch the KVM between the blades.

Fresh air plenums are at the top and bottom of the chassis. The bottom fresh air plenum provides non-preheated air to the M1000e power supplies. The top fresh air plenum provides non-preheated air to the CMC, iKVM and I/O modules.

Front control panel

The front control panel contains the local user interface. Functions include interactive graphical LCD, server and enclosure information and troubleshooting, and power button. This device is hot-pluggable and is always powered, even in chassis standby mode.

The LCD panel on the M1000e chassis shows extensive information about the status of each hardware module, network information for the CMC and each iDRAC with Lifecycle Controller, and status messages with detailed explanations in plain language. You can access a wide variety of information about modules using the panel, including your type, user-defined name, configurations,



service tag numbers, and IP address information. The LCD panel can be retracted into the chassis (Figure 3), or extended and angled once deployed for full visibility no matter where the M1000e is mounted in the rack (Figure 4).

Figure 3. M1000e LCD panel in recessed position



Figure 4. M1000e LCD panel during usage



The LCD panel can be used as a diagnostic source and as a place to configure parameters of certain chassis components as well as the server's iDRAC with Lifecycle Controller network configuration. The primary function of the LCD panel is to provide real-time information on the health and status of the modules in the enclosure. LCD panel features include:

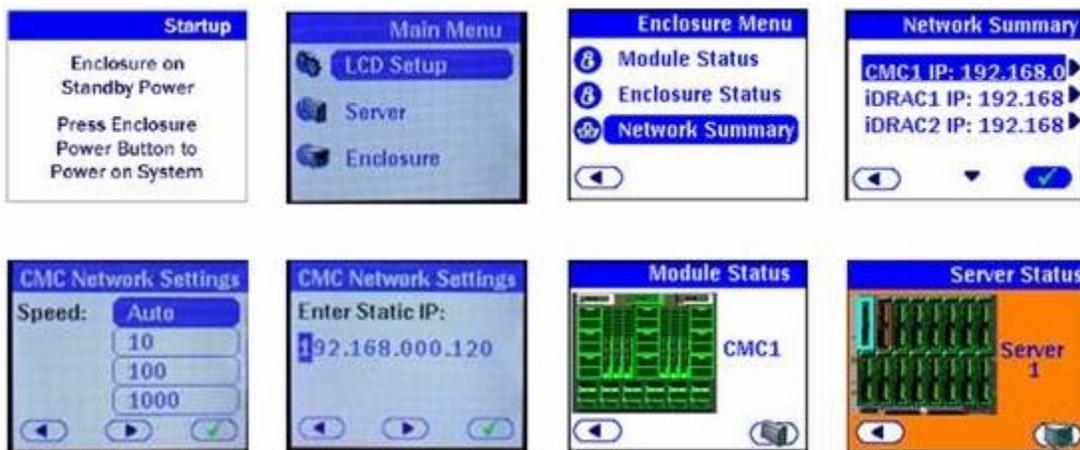
- A deployment setup wizard that allows you to configure the CMC module's network settings during initial system set up
- Menus to configure the iDRAC with Lifecycle Controller in each blade
- Status information screens for each blade



- Status information screens for the modules installed in the back of the enclosure, including the I/O modules, fans, CMC, iKVM, and power supplies
- A network summary screen listing the IP addresses of all components in the system
- Real time power consumption statistics, including high and low values and average power consumption
- Ambient temperature values
- AC power information
- Critical failure alerts and warnings

Figure 5 shows some of the capabilities of the LCD panel.

Figure 5. LCD panel capabilities



For more details on the capabilities of the LCD panel, see the following guides on Dell.com/Support/Manuals:

- *Dell PowerEdge M1000e, M915, M910, M820, M710HD, M710, M620, M610x, M610, M520, and M420 Getting Started Guide*
- *Dell Chassis Management Controller Firmware User's Guide*

Front ports

The M1000e supports one VGA port and two USB keyboard/mouse ports.

Back view and features

The back of the M1000e enclosure contains system management, cooling, power, and I/O components. At the top of the enclosure are slots for two CMC cards and one integrated KVM switch. The enclosure ships by default with a single CMC, with the option of adding a second CMC to provide a fully redundant, active-standby, fault-tolerant solution for management access and control.

Interleaved in the center of the chassis are fans and I/O modules. This arrangement optimizes the balance of airflow through the system, allowing lower pressure build-up in the system and resulting in lower airflow requirements for the fans.



Figure 6. Back of enclosure



Fans

The PowerEdge M1000e chassis comes standard with nine hot-plug, redundant fan modules that are distributed evenly across the back of the enclosure as shown in Figure 7. The speed of each fan is individually managed by the CMC. Together, these design innovations provide the following:

- Significant power savings as compared to older servers
- Less airflow required as compared to the same number of similarly configured 1U servers
- Similar sound-level as compared to previous servers

A similar acoustic profile as compared to previous servers, fans are N+1 redundant, meaning that any single fan can fail without impacting system uptime or reliability. In the event of a fan failure, system behavior is dependent on the resultant temperatures of the system, as monitored by the server module iDRAC with Lifecycle Controller and I/O modules. The CMC continues to interpret the airflow needs of each server and I/O module to control the fan speeds appropriately.

The system will not ramp the fans to full speed in the event of a fan failure unless deemed necessary by onboard monitoring. Failure of more than one fan will not automatically result in shutting down the blade servers because they have their own self-protection mechanisms to prevent them from running too hot. The result of a failure of multiple fans would depend on the configuration, ambient temperature, and workload being run. For example, the processors within a blade are automatically



throttled back by that server if they reach a thermal threshold and then shut down if a critical over-temperature threshold is met.

Note: The blank blade, hard drive, and server I/O fillers for every blank slot are required for cooling and airflow reasons.

Figure 7. Installed CMC, I/O module, and power supply blanks



I/O modules

For more information about the M1000e enclosure I/O modules, see Networking and Input/Output section of this guide.

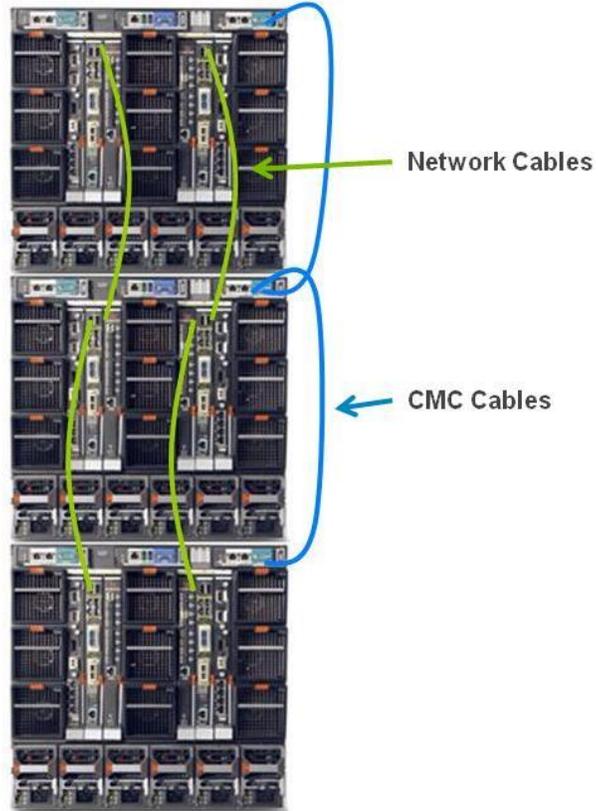
Cabling

The M1000e enclosure offers two types of external cabling simplification features:

- Stacked Ethernet Switching:
 - Internal switches have optional 10GbE uplinks and/or stacking connectors
 - Manage and configure multiple switches as one with stacking
 - Consolidate uplinks from multiple chassis into two to four 10GbE ports
- Stacked CMCs:
 - CMC has a second Ethernet port for connection to other CMCs in the rack
 - CMC connects to the management network to manage all blade servers
 - Saves port consumption on external switches



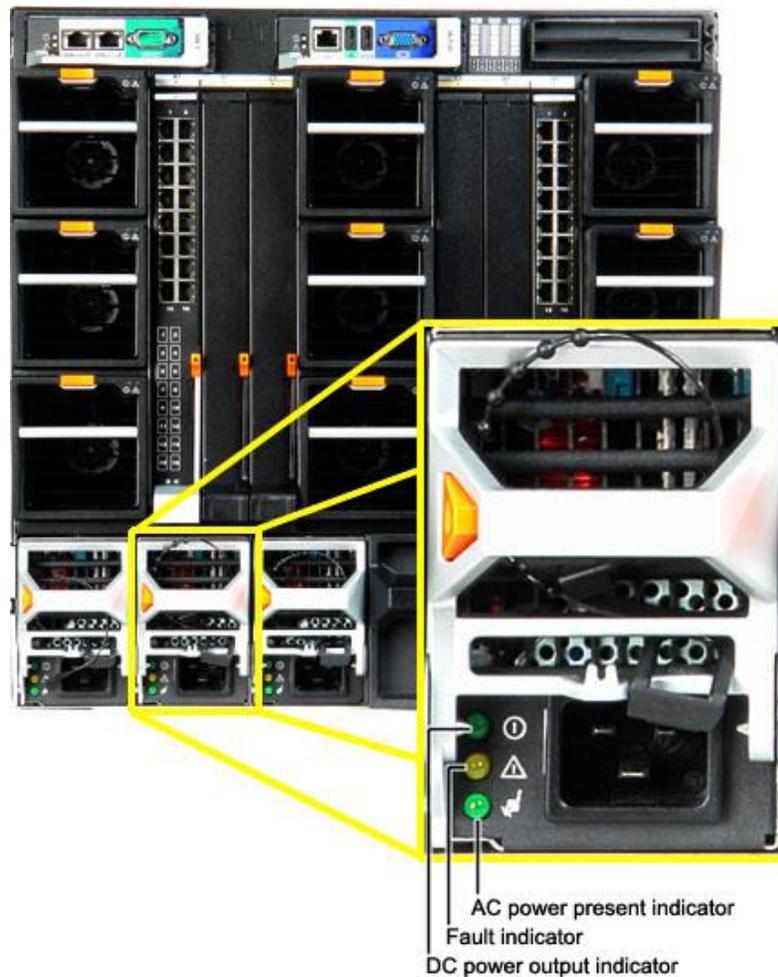
Figure 8. Simplified cabling



Power supply indicators

Figure 9 shows the power supply indicators.

Figure 9. Power supply indicators



Security features

The M1000e offers many security features, including the ability to:

- Assign one administrator per blade or one administrator per multiple blades
- Grant permissions to some blades but not to others
- Customize administrative access for CMC, iDRAC with Lifecycle Controller, and I/O

Most of the security capabilities are driven by the CMC, which provides centralized configuration of the M1000e enclosure's security settings and user access, and is secured by a user-modifiable password. The CMC's security features include:

- User authentication through optional Active Directory and LDAP services or hardware-stored user IDs and passwords
- Role-based authority, which enables an administrator to configure specific privileges for each user

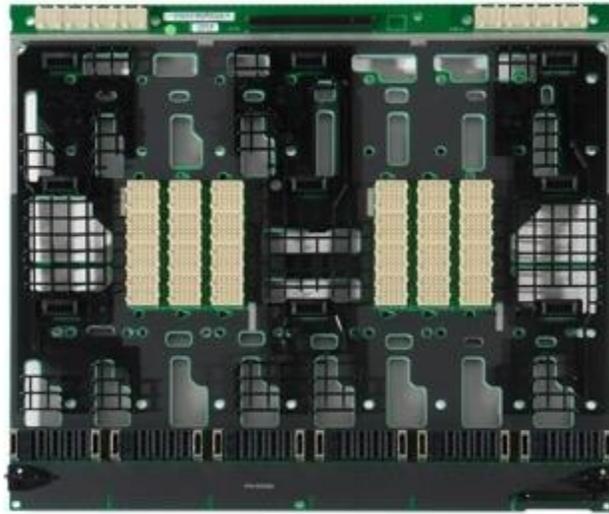
- User ID and password configuration through a Web interface that supports 128-bit SSL 3.0 encryption and 40-bit SSL 3.0 encryption (for countries where 128-bit is not acceptable)
- Configurable IP ports (where applicable)
- Login failure limits per IP address, with login blocking from the IP address when the limit is exceeded
- Configurable session auto time out and number of simultaneous sessions
- Limited IP address range for clients connecting to the CMC
- Secure Shell (SSH), which uses an encrypted layer for higher security
- Single Sign-on, Two-Factor Authentication and Public-Key Authentication
- Front panel access that can be disabled



4 Midplane

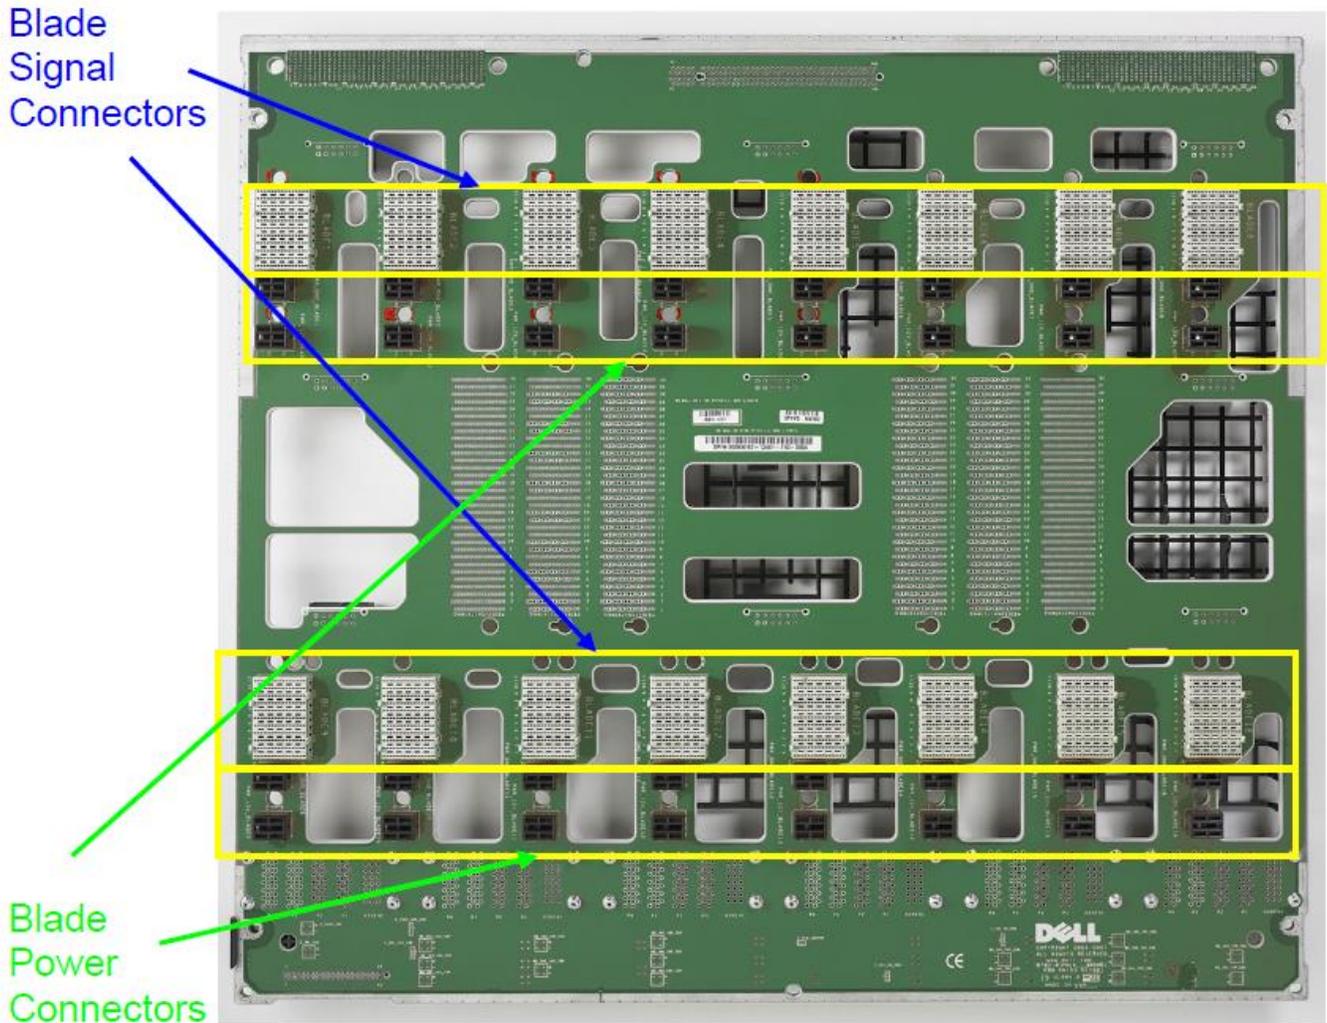
Though hidden from view in an actively running system, the midplane is the focal point for all connectivity within the M1000e modular system. The midplane is a large printed circuit board providing power distribution, fabric connectivity, and system management infrastructure. Additionally it allows airflow paths for the front-to-back cooling system through ventilation holes.

Figure 10. Midplane



As is requisite for fault-tolerant systems, the M1000e midplane is completely passive, with no hidden stacking midplanes or interposers with active components. I/O fabrics and system management are fully redundant from each hot-pluggable item. The system management Ethernet fabric is fully redundant when two CMCs are installed, with two point-to-point connections from each server module.

Figure 11. M1000e midplane front view

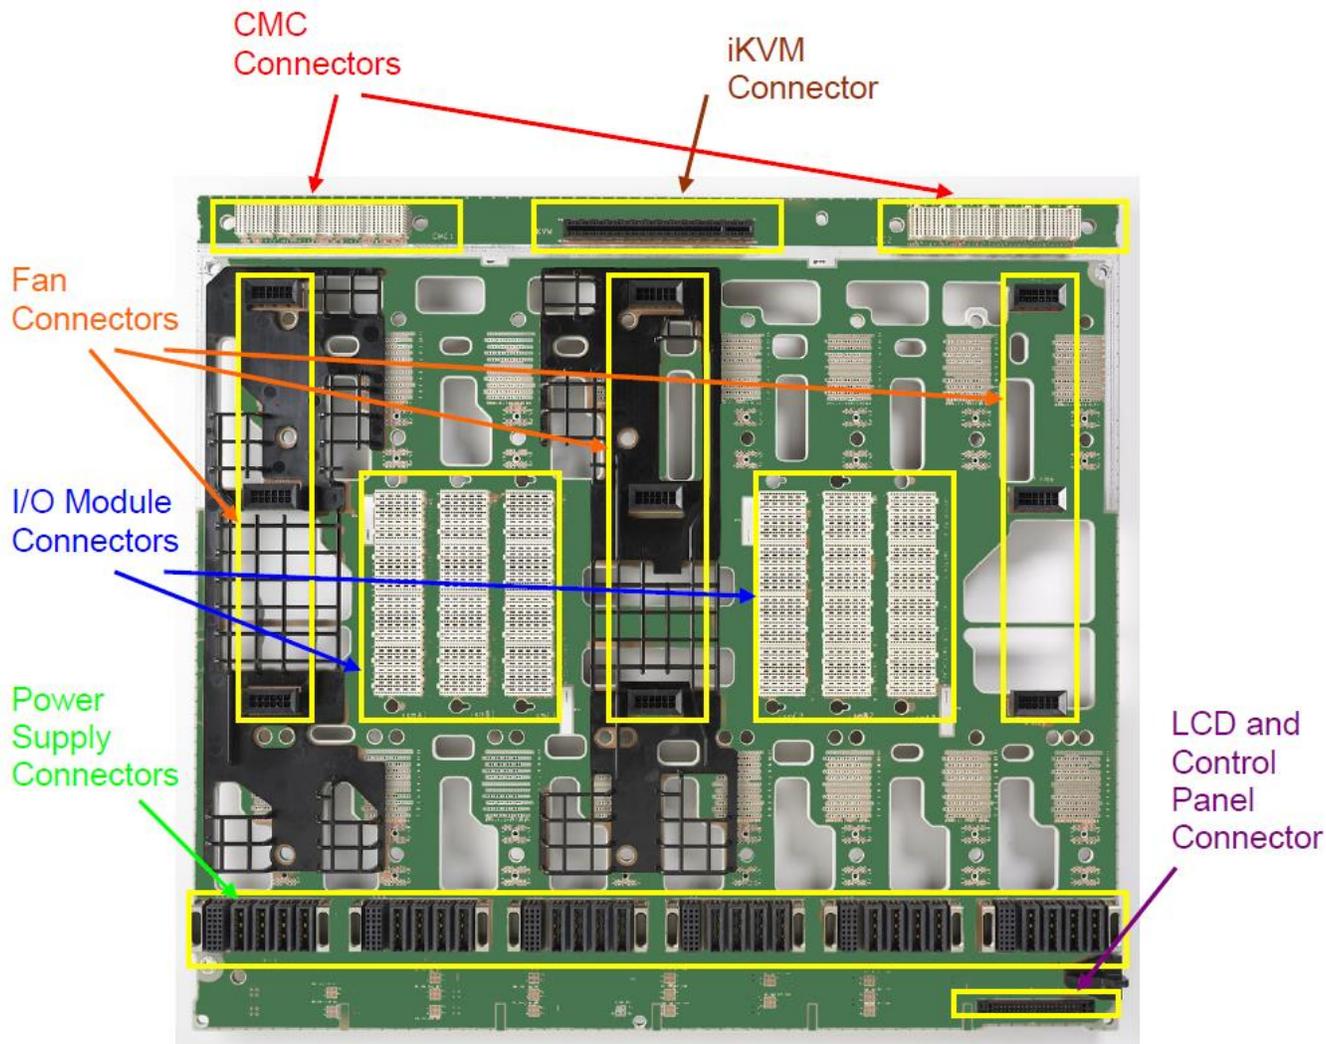


The system is designed for receptacles on all midplane connectors and pins on all pluggable components, so any potential for bent pins is limited to the pluggable field replaceable unit, not to the system. This contributes to the high reliability and uptime of the M1000e modular system.

The midplane is physically attached to the enclosure front structural element. It is aligned by guide-pins and edges in all three axes. This provides close tolerance alignment between the server modules and their midplane connections.



Figure 12. M1000e midplane back view



All M1000e midplane routing is fully isolated, supporting all chassis power, fabric, system management and fault-tolerance requirements.



5 Processors and Memory

See the *Dell PowerEdge Mxxx Technical Guide* on Dell.com/PowerEdge for each of the compatible blade servers offered for more details on processors and memory offered.



6 Storage

The M1000e is designed to provide the flexibility to support various storage solutions depending on customer requirements. With the latest addition of the Dell EqualLogic PS-M4110 to the blade ecosystem, customers now have greater flexibility from a shared storage perspective within the M1000e chassis.

The EqualLogic PS-M4110 Blade Array is the first generation iSCSI storage blade for the EqualLogic product line and offers features consistent with the PS4100 series. The PS-M4110 fits inside the Dell PowerEdge M1000e blade chassis and gives PowerEdge M-Series blade server customers shared storage inside the chassis. For more information about EqualLogic PS-M4110 Blade Array, visit Dell.com/EqualLogic.



7 Networking and Input/Output

Embedded NICs/LAN on Motherboard

For information on embedded NICs/LAN on Motherboard (LOM), see the *Dell PowerEdge Mxxx Technical Guide* on Dell.com/PowerEdge for each of the compatible blade servers.

Input/Output

The Dell PowerEdge M-series servers provide complete, snap-in FLEX™ I/O scalability down to the switch interconnects. FLEX I/O technology is the foundation of the M1000e I/O subsystem. You can mix and match I/O modules from a wide variety of options including Cisco, Dell PowerConnect, Fibre Channel, and InfiniBand options. The I/O modules may be installed singly or in redundant pairs. For more information, see the *Dell PowerEdge M1000e Enclosure Owner's Manual* on Dell.com/Support/Manuals.

I/O modules connect to the blades through three redundant I/O fabrics. The enclosure was designed for 5+ years of I/O bandwidth and technology.

The I/O system offers a wide variety of options to meet nearly any network need:

- Complete, on-demand switch design
- Easily scale to provide additional uplink and stacking functionality
- No need to waste your current investment with a rip-and-replace upgrade
- Flexibility to scale Ethernet stacking and throughput
- Partnered Solutions with Cisco, Emulex, and Brocade®
- QDR, FDR10 and FDR InfiniBand switch options available for HPCC
- Up to 8 high-speed ports
- Cisco virtual blade switch capability
- Ethernet port aggregator
- Virtualization of Ethernet ports for integration into any Ethernet fabric
- Fibre Channel products from Brocade and Emulex offering powerful connectivity to Dell Compellent SAN fabrics
- High-availability clustering inside a single enclosure or between two enclosures (for details go to Dell.com/Support/Manuals for documentation on clustering, organized by type of storage)

Each server module connects to traditional network topologies while providing sufficient bandwidth for multi-generational product lifecycle upgrades. I/O fabric integration encompasses networking, storage, and interprocessor communications (IPC).

Definitions:

- A **fabric** is defined as a method of encoding, transporting, and synchronizing data between devices. Examples of fabrics are Gigabit Ethernet (GbE), Fibre Channel (FC) or InfiniBand (IB). Fabrics are carried inside the M1000e system, between server module and I/O modules through the midplane. They are also carried to the outside world through the physical copper or optical interfaces on the I/O modules.
- A **lane** is defined as a single fabric data transport path between I/O end devices. In modern high speed serial interfaces each lane is comprised of one transmit and one receive differential pair. In reality, a single lane is four wires in a cable or traces of copper on a printed circuit board, a



transmit positive signal, a transmit negative signal, a receive positive signal, and a receive negative signal. Differential pair signaling provides improved noise immunity for these high speed lanes. Various terminology is used by fabric standards when referring to lanes. PCIe calls this a lane, InfiniBand calls it a physical lane, and Fibre Channel and Ethernet call it a link.

- A **link** is defined here as a collection of multiple fabric lanes used to form a single communication transport path between I/O end devices. Examples are two-, four- and eight-lane PCIe, or four-lane 10GBASE-KX4. PCIe, InfiniBand and Ethernet call this a link. The differentiation has been made here between lane and link to prevent confusion over Ethernet's use of the term link for both single and multiple lane fabric transports. Some fabrics such as Fibre Channel do not define links as they simply run multiple lanes as individual transports for increased bandwidth. A link as defined here provides synchronization across the multiple lanes, so they effectively act together as a single transport.
- A **port** is defined as the physical I/O end interface of a device to a link. A port can have single or multiple lanes of fabric I/O connected to it.

Quantities and priorities

There are three supported high-speed fabrics per M1000e half-height server module, with two flexible fabrics using optional plug-in mezzanine cards on the server, and one connected to the LOMs on the server. The ports on the server module connect through the midplane to the associated I/O modules (IOM) in the back of the enclosure, which then connect to LAN/SAN/IPC networks.

The optional mezzanine cards are designed to connect through the eight-lane PCIe to the server module's chipset in most cases. Mezzanine cards have either one dual port ASIC with four- or eight-lane PCIe interfaces or dual ASICs, each with four-lane PCIe interfaces. External fabrics are routed through high-speed, 10Gbps-capable air dielectric connector pins through the planar and midplane. For best signal integrity, the signals isolate transmit and receive signals for minimum crosstalk. Differential pairs are isolated with ground pins and signal connector columns are staggered to minimize signal coupling.

The M1000e system management hardware and software includes fabric consistency checking, preventing accidental activation of any misconfigured fabric device on a server module. The system will automatically detect a misconfiguration and alert you of the error. No damage occurs to the system, and you can reconfigure the faulted module.

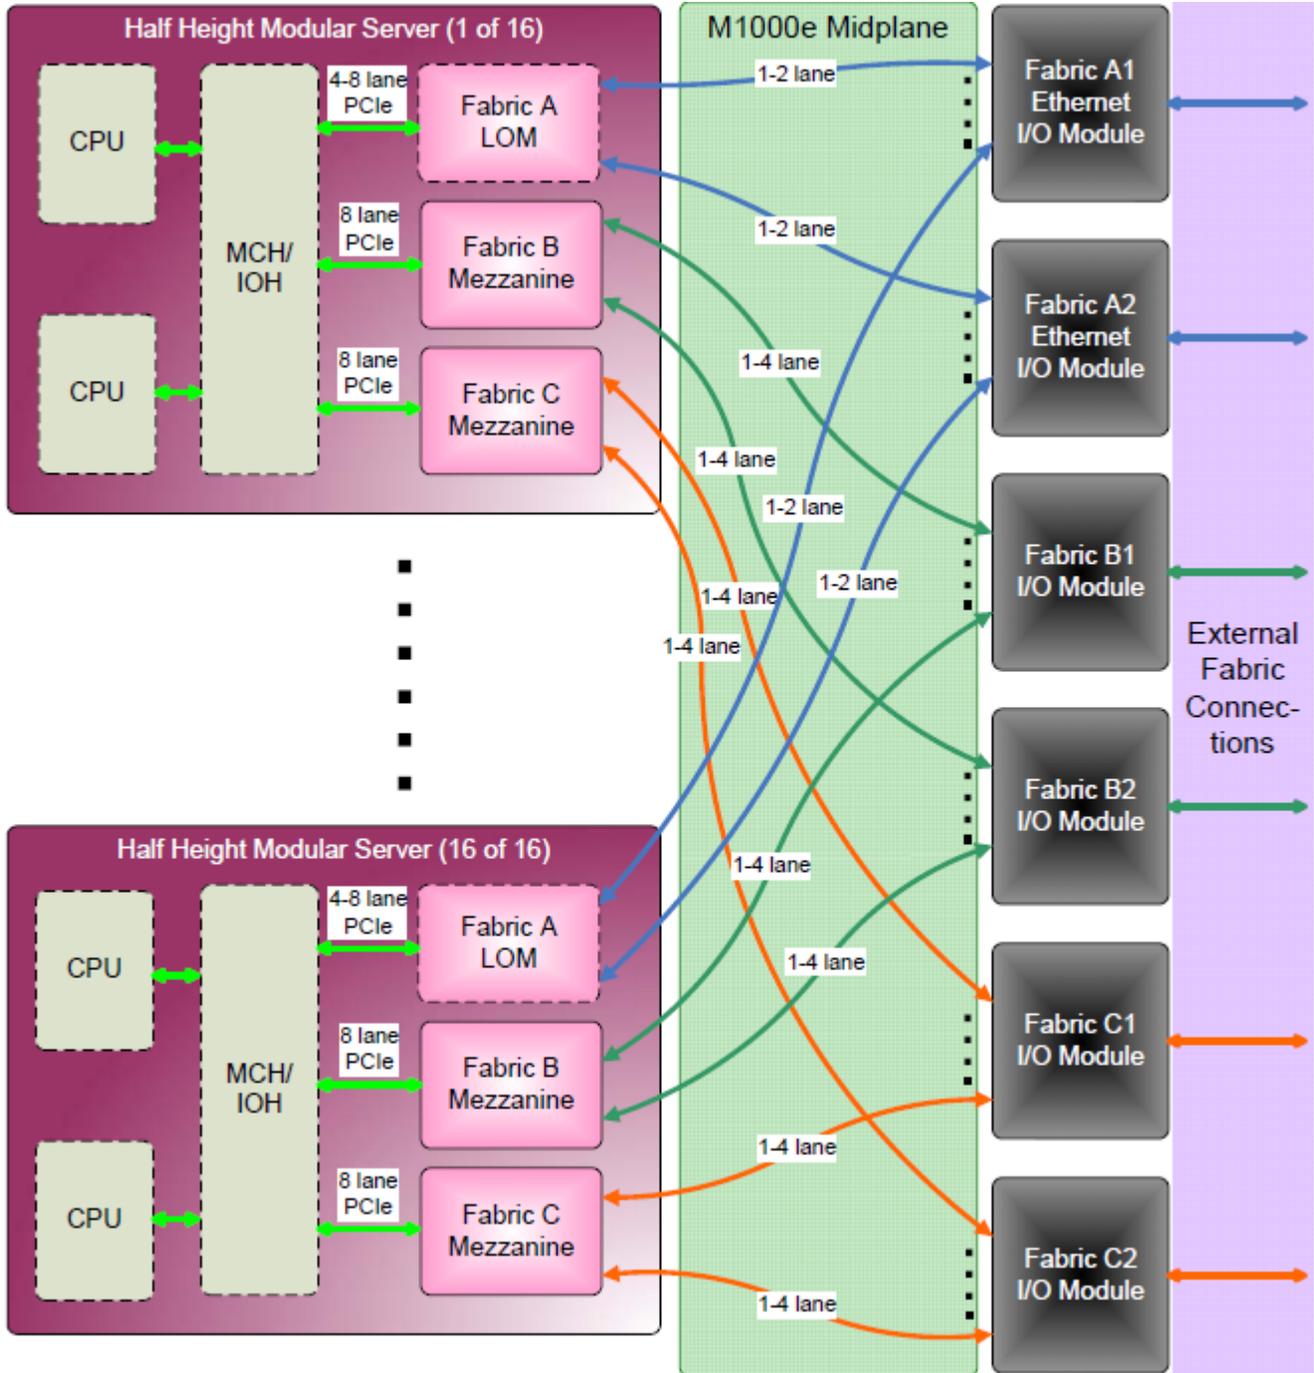
M1000e I/O is fully scalable to current and future generations of server modules and I/O Modules. There are three redundant multi-lane fabrics in the system, as illustrated in Figure 13.

In its original configuration, the M1000e midplane is enabled to support up to four Gigabit Ethernet links per server module on Fabric A. Thus, potential data bandwidth for Fabric A is 4 Gbps per half-height server module. A future midplane upgrade may enable higher bandwidth on Fabric A.

The M1000e provides full 10/100/1000M Ethernet support when using Ethernet pass-through modules enabling you to connect to any legacy infrastructure whether using Ethernet pass-through or switch technology. This technical advance uses in-band signaling on 1000Base-KX transport and requires no user interaction for enablement.



Figure 13. High-speed I/O architecture



Fabric B and C are identical, fully customizable fabrics, routed as two sets of four lanes from mezzanine cards on the server modules to the I/O modules in the back of the chassis. Supported bandwidth ranges from 1 to 10 Gbps per lane depending on the fabric type used.



Table 2. Fabric specifications

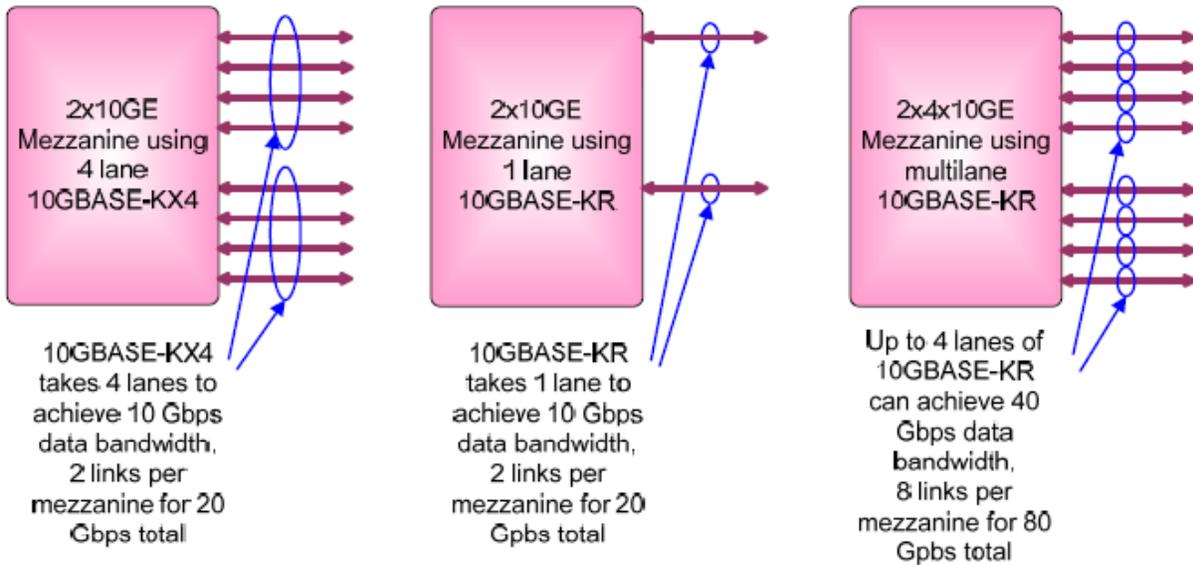
Fabric	Encoding	Symbol rate per lane (Gbps)	Data rate per lane (Gbps)	Data rate per link (Gbps)	Lanes per link per industry specification
PCIe 1.0	8B/10B	2.5	2	8 (4 lane)	1, 2, 4, 8, 12, 16, 32
PCIe 2.0	8B/10B	5	4	16 (4 lane)	1, 2, 4, 8, 12, 16, 32
PCIe 3.0	128B/130B	8	7.88	31.51 (4 lane)	1, 2, 4, 8, 12, 16, 32
SATA 3Gbps	8B/10B	3	2.4	2.4	1
SATA 6Gbps	8B/10B	6	4.8	4.8	1
SAS 3Gbps	8B/10B	3	2.4	2.4	1-Any
SAS 6Gbps	8B/10B	6	4.8	4.8	1-Any
FC 4Gbps	8B/10B	4.25	3.4	3.4	1
FC 8Gbps	8B/10B	8.5	6.8	6.8	1
FC 16Gbps	64B/66B	14.025	13.6	13.6	1
IB SDR	8B/10B	2.5	2	8 (4 lane)	4, 12
IB DDR	8B/10B	5	4	16 (4 lane)	4, 12
IB QDR	8B/10B	10	8	32 (4 lane)	4, 12
IB FDR10	64B/66B	10.3125	10	40 (4 lane)	4, 12
IB FDR	64B/66B	14.0625	13.64	54.56 (4 lane)	4, 12
GbE: 1000BASE-KX	8B/10B	1.25	1	1	1
10GbE: 10GBASE-KX4	8B/10B	3.125	2.5	10 (4 lane)	4
10GbE: 10GBASE-KR	64B/66B	10.3125	10	10	1

Symbol rate is defined as the maximum fabric speed with all data encoded as per its encoding mechanism. Data rate is the maximum amount of actual data bits that can be transported on the fabric. Some fabrics are defined as single lane only, while others are defined in various lane increments for link to lane aggregation in order to provide increased total bandwidth.

The bandwidth growth path for Ethernet is from GE to 10GE, supported through three backplane standards defined by the IEEE 802.3ap Committee. 1000BASE-KX supports GE through 1.25 Gbps differential backplane signaling. 10GBASE-KX4 supports 10GE through four lanes of 3.125 Gbps differential backplane signaling. The four lanes are aggregated to provide the sum total 10 Gbps bandwidth requirement. 10GBASE-KR is an emerging IEEE 802.3ap standard that supports 10GE through a single lane of 10.3125 Gbps differential backplane signaling. In terms of customer impact, this means that today's modular systems using 10GBASE-KX4 are limited to one 10GE channel per four lanes of internal differential signaling, while in the future those same lanes can carry multiple lanes of 10GBASE-KR and provide a resultant increase in total system throughput. This growth path applies to both Fabrics B and C in the M1000e.



Figure 14. Ethernet growth path



GbE fabrics use single lanes per link over 1000BASE-KX. Initially 10GbE will be supported through 10GBASE-KX4, which uses four lanes per link, as does InfiniBand. As technology matures and enables high level integration, 10GBASE-KR will replace 10GBASE-KX4 as the 10GbE backplane transport of choice, enabling 10 Gbps over single lanes within the system midplane. In the future, multiple lanes of 10GBASE-R per mezzanine will allow expanded Ethernet bandwidth up to 80 Gbps per fabric mezzanine.

Assuming a full population of GE and 10GbE lanes in Fabrics A, B and C, the M1000e can deliver theoretical total bandwidth of 8.4 Terabits per second.

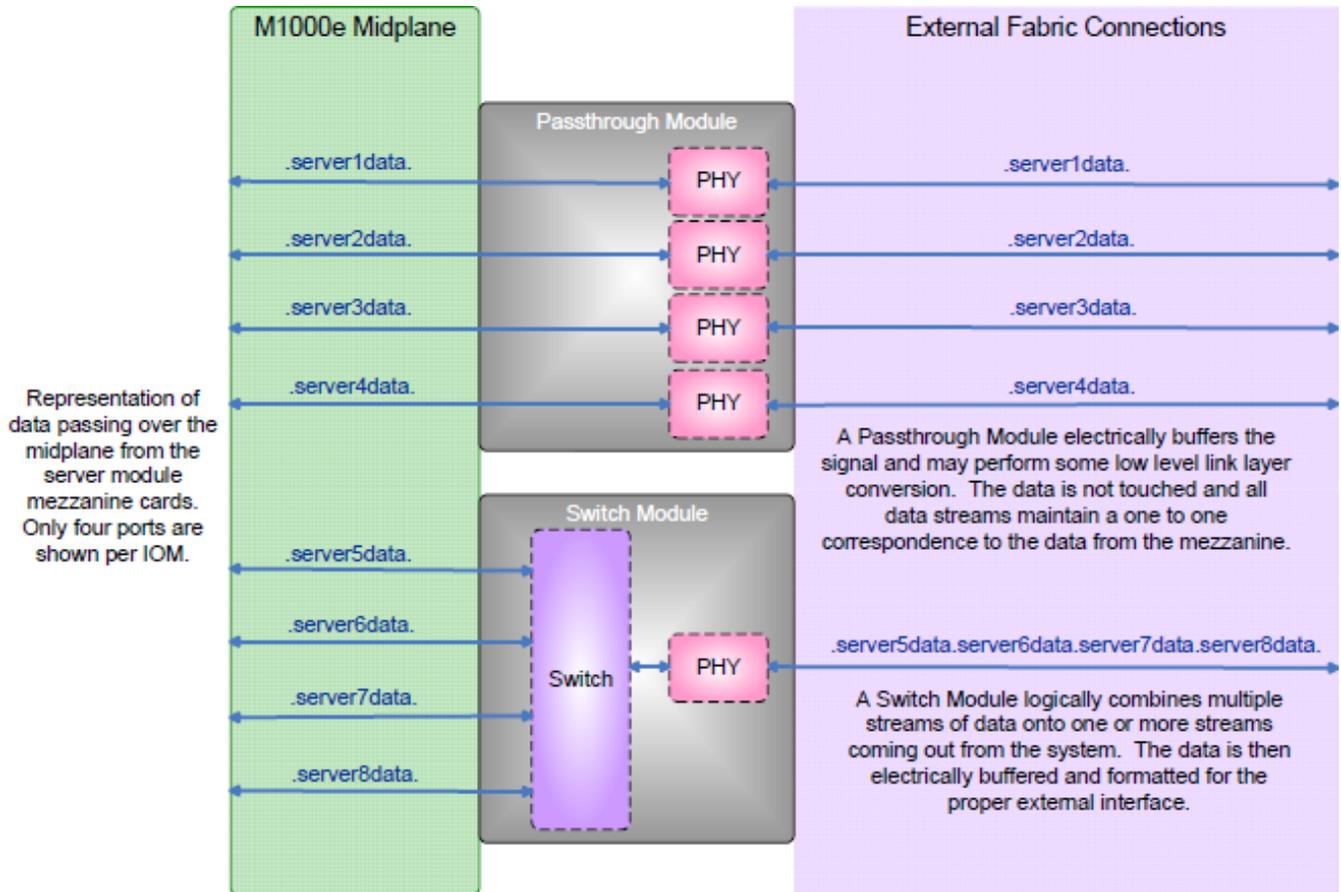
The M1000e is designed for full support of all near-, medium-, and long-term I/O infrastructure needs. While the M1000e system's bandwidth capabilities lead the industry, the M1000e is also intelligently designed for maximum cost, flexibility, and performance benefit.

While Fabric A is dedicated to the server module LOMs, requiring Ethernet switch or pass-through modules for I/O slots A1 and A2, Fabrics B and C can be populated with Ethernet, Fibre Channel, or InfiniBand solutions.

I/O modules are used as pairs, with two modules servicing each server module fabric providing a fully redundant solution. I/O modules may be pass-throughs or switches. Pass-through modules provide direct 1:1 connectivity from each LOM/mezzanine card port on each server module to the external network. Switches provide an efficient way to consolidate links from the LOM or mezzanine cards on the server modules to uplinks into the network.



Figure 15. Difference between pass-through and switch modules



Supported mezzanine cards and switches

Dell supports one mezzanine design standard and one I/O module design standard for true modular computing. Supported I/O modules include:

- PowerConnect M6220 switch (GbE, 10GbE uplinks)
- PowerConnect M6348 switch (high port count 1GbE, 10GbE uplinks)
- PowerConnect M8024-k (10Gb Ethernet)
- Dell Networking MXL 10/40GE
- Dell PowerEdge M I/O Aggregator
- Dell M8428-k (10Gb converged networking/FCoE)
- Cisco Catalyst 3032 (1Gb Ethernet)
- Cisco Catalyst 3130G (1Gb Ethernet, stacking)
- Cisco Catalyst 3130x (1GbE+ 10GbE uplinks, stacking)
- 1Gb Ethernet pass-through module
- 10Gb Ethernet pass-through module (original model, for 10Gb XAUI mezz cards)
- 10Gb Ethernet pass-through II (for 10Gb XAUI mezz cards)
- 10Gb Ethernet pass-through-k (for 10Gb KR mezz cards/NDCs)
- Brocade M5424 (8/4Gbps Fibre Channel switch)
- Brocade M4424 (4Gbps Fibre Channel switch)

- Dell 8/4Gbps FC SAN module (NPIV-based FC aggregator)
- Dell 8/4Gbps FC pass-through module
- Emulex 4Gbps FC pass-through module
- Mellanox M4001Q (QDR/40Gb InfiniBand switch; single-wide switch module)
- Mellanox M4001T (FDR/40Gb InfiniBand switch; single-wide switch module)
- Mellanox M4001F (FDR/56Gb InfiniBand switch; single-wide switch module)

For an up-to-date list of supported I/O modules, see the [Ethernet Mezzanine Cards](#) site on [Dell.com](#).

I/O module installation

For detailed information on installing the I/O modules in your system, see the *Dell PowerEdge M1000e Enclosure Owner's Manual* on [Dell.com/Support/Manuals](#).

FlexAddress

FlexAddress™ delivers persistent storage and network identities, equipping a data center to handle predictable or even unplanned changes—increase, upgrade, or replace servers without affecting the network or storage and minimizing downtime.

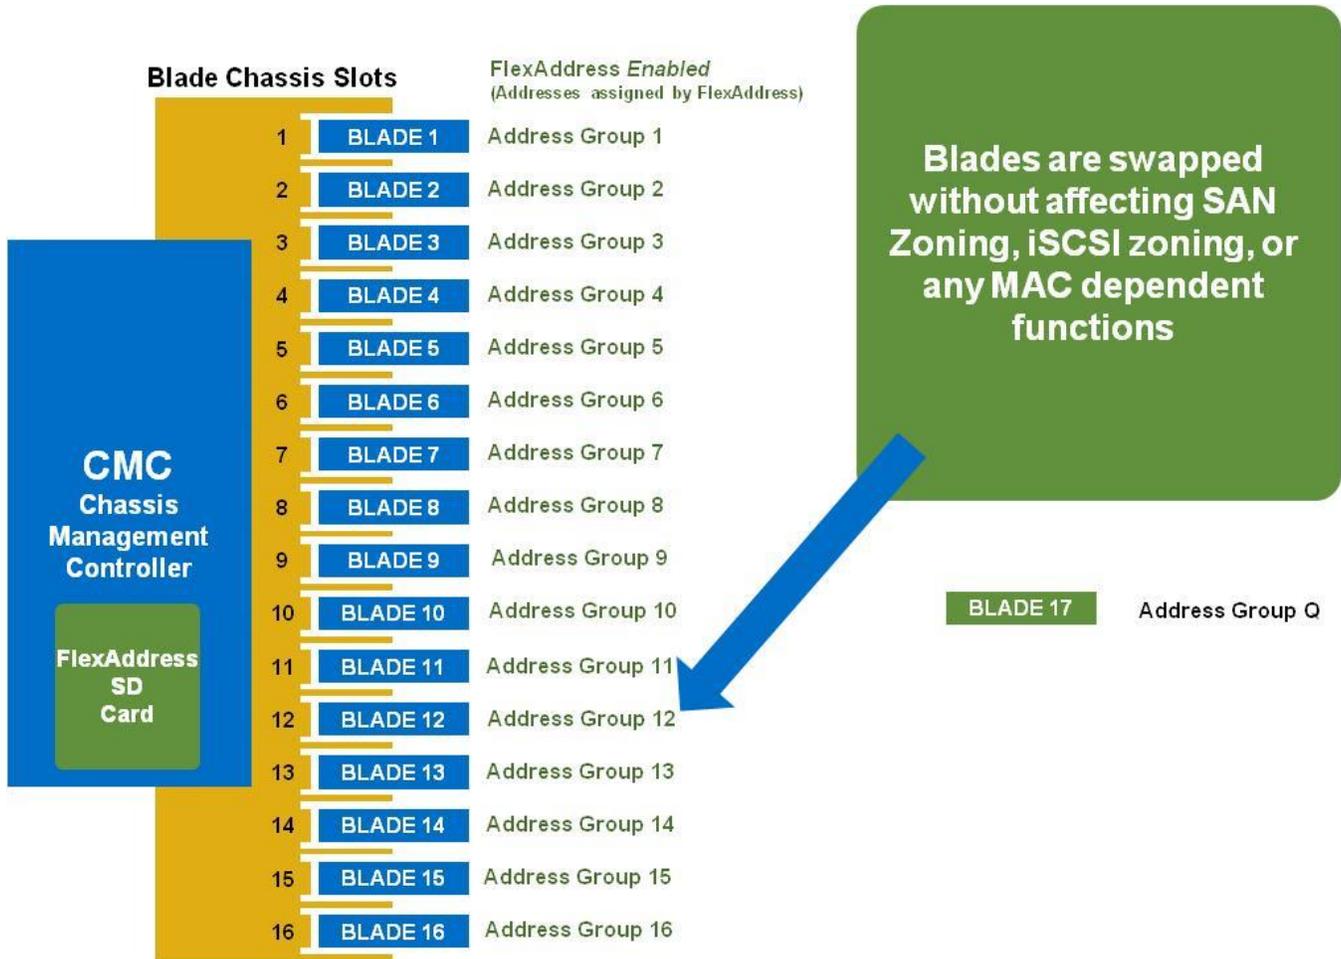
Dell's patent-pending FlexAddress technology allows any M-Series blade enclosure to lock the World Wide Name (WWN) of the Fibre Channel controller and Media Access Control (MAC) of the Ethernet and iSCSI controller into a blade slot, instead of to the blade's hardware. By removing the network and storage identity from the server hardware, you can now upgrade and replace components or the entire server without changing the identity on the network. This technology works with any vendor's installed I/O module as well as with Dell PowerConnect products.

FlexAddress delivers the ability to:

- Service a blade or I/O mezzanine card, upgrade the I/O mezzanine cards to newer technology, or upgrade the entire server with new technology while maintaining the mapping to Ethernet and storage fabrics. This capability allows quick, painless connection and reduces downtime. This capability is especially powerful when operating in a boot-from-SAN environment.
- Quickly obtain a list of all MAC/WWNs in the chassis by slot and be assured these will never change.
- Efficiently integrate into existing management and network infrastructure.



Figure 16. FlexAddress addresses



FlexAddress replaces the factory-assigned World Wide Name/Media Access Control (WWN/MAC) IDs on a blade with WWN/MAC IDs from the FlexAddress SD card associated with that slot. This user-configurable feature enables a choice of iSCSI MAC, Ethernet MAC, and/or WWN persistence enabling blades to be swapped without affecting SAN zoning, iSCSI zoning, or any MAC-dependent functions. The write-protected FlexAddress SD card comes provisioned with a unique pool of 208 MACs and 64 WWNs. Other types of SD cards inserted into the CMC's SD card slot are ignored.

You can order FlexAddress with a new enclosure, or implement FlexAddress for an existing enclosure by purchasing a customer kit. If you purchase FlexAddress with an enclosure, it will be active when the system is powered up. For an existing enclosure, you will need to add a FlexAddress SD card to the CMC, and upgrade the iDRAC with Lifecycle Controller firmware, Ethernet and Fibre Channel controllers' firmware, server BIOS, and CMC firmware. All blades and CMC must have the correct versions of firmware to properly support this feature. If you install redundant CMCs, it is not necessary to add an SD card to both CMCs since the WWN/MAC addresses are pushed to the chassis control panel when enabled for redundancy. If one CMC becomes inoperable, the other CMC still has access to the WWN/MAC addresses in the control panel. Blades that are up and running are not affected since they already have their WWN/MACs programmed into their controllers. If you replace the control panel, the SD card will push the WWN/MACs back to the replacement control panel. Additionally, the chassis control panel also stores CMC configuration information, so keeping a backup of the CMC configuration file it recommended.



The CMC manages the following functions specific to FlexAddress:

- Provides a user interface for enabling or disabling the FlexAddress feature — on a per-blade-slot basis, a per-fabric basis, or both
- Identifies and reports device information for each of the supported fabric types — LOMs, Ethernet and Fibre Channel mezzanine cards
- Validates all the components of the FlexAddress feature — SD card validation, System BIOS, I/O controller firmware, CMC firmware and iDRAC firmware versions
- Reports FlexAddress feature status for each of the blade slots
- Provides information for both server-assigned (factory-programmed) and chassis-assigned (FlexAddress) addresses on each supported device
- Logs any system-level errors that may prevent the FlexAddress feature from being used on the chassis or on a given device

Figure 17. FlexAddress screen in the CMC

Slot 15 FlexAddress is enabled			
Location	Fabric	Server-Assigned	Chassis-Assigned
A1	Gigabit Ethernet	00:1D:09:FC:BD:34	✓00:1D:09:FF:BB:37
	iSCSI	00:1D:09:FC:BD:35	✓00:1D:09:FF:BB:38
A2	Gigabit Ethernet	00:1D:09:FC:BD:36	✓00:1D:09:FF:BB:39
	iSCSI	00:1D:09:FC:BD:37	✓00:1D:09:FF:BB:3A
B1	None		
B2	None		
C1	None		
C2	None		

Table 3. FlexAddress features and benefits

Features	Benefits
Locking the WWN of the Fibre Channel controller and the MAC of the Ethernet and iSCSI controller into a blade slot instead of to the blade's hardware	Easily replace blades without network management effort
Servicing or replacing a blade or I/O mezzanine card and maintain all address mapping to Ethernet and storage fabrics	Ease-of-management
Easy and highly reliable booting from Ethernet or Fibre Channel based Storage Area Networks (SANs)	An almost no-touch blade replacement
All MAC/WWN/iSCSIs in the chassis will never change	Fewer future address name headaches
Fast and efficient integration into existing network infrastructure	No need to learn a new management tool Low cost versus switch-based solution
FlexAddress: simple and easy to implement	Simple and quick to deploy
FlexAddress SD card with a unique pool of MAC/WWNs — can be enabled on a single enclosure at a given time, until disabled	No need for the user to configure No risk of duplicates on your network or SAN
Works with all I/O modules including Cisco, Brocade, and Dell PowerConnect switches as well as pass-through modules	Choice being independent of switch or pass-through module



Wake-on-LAN (WOL) is enabled after a power-down and power-up of the enclosure. FlexAddress on Ethernet devices is programmed by the module server BIOS. In order for the blade BIOS to program the address, it must be operational, which requires the blade to be powered up. Once the power-down and power-up sequence completes, the FlexAddress feature is available for WOL function.

For information on implementing FlexAddress on an M1000e, see the *Dell Chassis Management Controller Firmware User's Guide* on Dell.com/Support/Manuals.



8 Power, thermal and acoustics

Lower overall system-level power draw is a result of Dell's breakthrough system design. The PowerEdge M1000e chassis enclosure and blade servers maximize performance-per-watt through a combination of power and cooling, energy efficient technologies, and tools. Additionally, the PowerEdge modular blade servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Built on Dell Energy Smart technology, the M1000e is one of the most power-efficient blade solutions on the market. The M1000e enclosure takes advantage of Energy Smart thermal design efficiencies, such as ultra-efficient power supplies and dynamic power-efficient fans with optimized airflow design to efficiently cool the chassis and enable better performance in a lower power envelope.

A modular system has many advantages over standard rack-mount servers in terms of power optimization, and this aspect was a focal point throughout the M1000e's conceptualization and development. The key areas of interest are power delivery and power management.

The M1000e provides industry-leading power efficiency and density, accomplished through highly efficient components, improved design techniques, and a fresh air plenum that reduces the air temperature to the power supply components. Lower operating temperature equates to higher power density for the power supply (exceeding 21 watts per cubic inch) and higher power efficiency (better than 87% at 20% load and higher at heavier loads, approaching 91% efficiency under normal operating conditions).

Power efficiency in the M1000e does not stop with the power supply. Every aspect of efficiency has been tweaked and improved from previous designs — adding more copper to PC board power planes to reduce I²R losses, improving inductors and other components, and replacing some linear voltage regulators with more efficient switching regulators.

Power supplies

The power distribution inside the M1000e modular server system consists of a power supply system located in the back of the chassis as shown in Figure 18.

Figure 18. Power supplies in M1000e



Figure 19. M1000e power supply



Dell power supplies use output ORing FETs to isolate the power supply from the 12V system bus. If a single power supply fails its output ORing FET, the power supply will turn off removing itself from the bus like an electrical switch that turns off when the power supply fails.

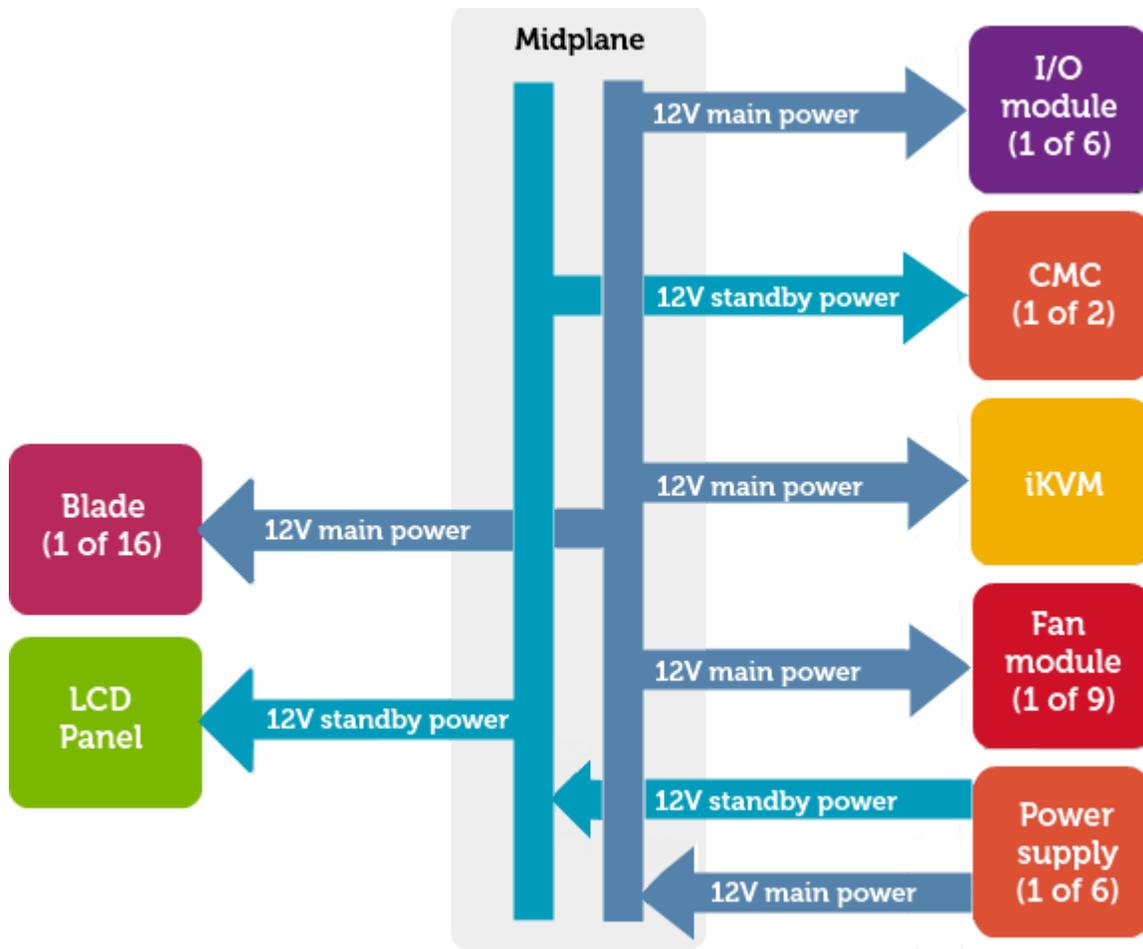
When DPSE is enabled, the PSU units move between on and off states depending on actual power draw conditions to achieve high power efficiency.

In the N+N power supply configuration, a system provides protection against AC grid loss or power supply failures. If one power grid fails, three power supplies lose their AC source, and the three power supplies on the other grid remain powered, providing sufficient power for the system to continue running. In the N+1 configuration, only power supply failures are protected, not grid failures. The likelihood of multiple power supplies failing at the same time is remote. In the N+0 configuration, there is no power protection and any protection must be provided at the node or chassis level. Typically this case is an HPCC or other clustered environment where redundant power is not a concern, since the parallelism of the processing nodes across multiple system chassis provides all the redundancy that is necessary.

The midplane carries all 12V DC power for the system, both main power and standby power. The CMCs, LCD and control panel are powered solely by 12V standby power, ensuring that chassis level management is operational in the chassis standby state, whenever AC power is present. The server modules, I/O modules, fans and iKVM are powered solely by 12V main power.



Figure 20. Power architecture



These supplies are hot-pluggable, so at any time you can move a power supply from one slot to another as long as there is enough reserve capacity.

2360W power supply

The 2360W (230V) power supply offers current sharing between power supplies, allowing a total system redundant power of approximately 6700W in a 3+3 power supply configuration.

The M1000e requires three power supplies to power a fully populated system or six power supplies in a fully redundant system.

Table 4. 2360W power supply

Input voltage	Input frequency	Output load	Power factor (PF)
230VAC	50-60 Hz	10%	.91
230VAC	50-60 Hz	20%	.97
230VAC	50-60 Hz	50%	.98
230VAC	50-60 Hz	100%	.99



2700W power supply

The 2700W power supply provides two different capacity levels: 2700W and 1350W. The capacity is determined by the AC cord line. If 220VAC is connected, the capacity is 2700W; if 110VAC is connected, the capacity is 1350W. A 220VAC and 110VAC line cannot be plugged into the chassis at the same time.

Table 5. 2700W power supply

Input voltage	Input frequency	Output load	Power factor (PF)
230VAC	50-60 Hz	10%	.94
230VAC	50-60 Hz	20%	.97
230VAC	50-60 Hz	50%	.99
230VAC	50-60 Hz	100%	.99

Power supply specifications

Each power supply offers:

- Up to 91%+ AC/DC conversion efficiency (2360W power supply) or up to 94%+ AC/DC conversion efficiency (2700W power supply)
- DPSE, which automatically engages the minimum number of supplies required to power a given configuration, maximizing power supply efficiency

The following detail the PowerEdge M1000e chassis power supply capabilities:

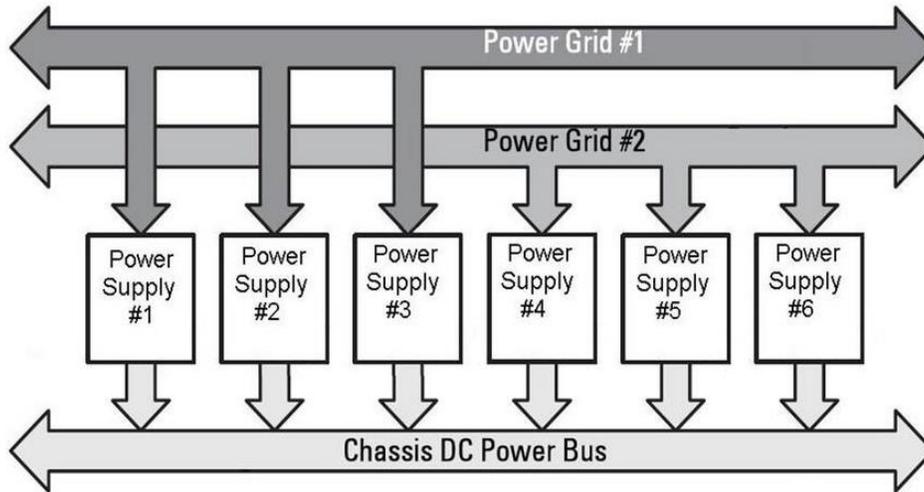
- 2360W or 2700W maximum for each power supply (depending which PSU is chosen)
- 220VAC (single PSU runs 180V–260VAC) or 110VAC input (2700W power supply only)
- 50Hz or 60Hz input
- 14A minimum (2360W power supply) to 18A maximum input (2700W power supply, running at 1350W during 100V operation)
- 192A (Amps) @ +12VDC output (operational)
- 4.5A @ +12VDC output (standby)

Power redundancy options

The six power supplies allow for two different redundancy options: grid redundancy and power supply redundancy. With grid redundancy, the six power supplies are divided into two banks. Bank 1 has slots 1 to 3, and bank 2 has slots 4 to 6. They should be populated in the following order: 1,4,2,5,3,6. See Figure 21 for grid redundancy setup.



Figure 21. Grid redundancy



With power supply redundancy, the six power supplies are seen as one bank. They should be populated in the following order: 1,2,3,4,5,6.

For information on the grid redundancy option, see the *Dell Chassis Management Controller Firmware User's Guide* on Dell.com/Support/Manuals.

Supported redundancy options in CMC 3.0

Table 6 lists and the following sections describe the power supply redundancy options available in CMC 3.0 for 1350W, 2700W, and 2360W power supplies.

Table 6. Supported power supply redundancy options

	1350W	2700W	2360W
No redundancy	1+0	1+0	1+0
	2+0	2+0	2+0
	3+0	3+0	3+0
	4+0	4+0	4+0
	5+0		5+0
	6+0		
AC redundancy	1+1	1+1	1+1
	2+2	2+2	2+2
	3+3	3+3	3+3
Power supply redundancy	1+1	1+1	1+1
	2+1	2+1	2+1
	3+1	3+1	3+1
	4+1	4+1	4+1
	5+1		5+1

Power tables for redundancy options

Table 7, Table 8 and Table 9 list the maximum power available for various power redundancy configurations.



Table 7. Grid redundancy

Configuration	Maximum power available (Watts DC)		
	2700W PSU	1350W PSU	2360W PSU
1+1	2534	1230	2222
2+2	5068	2460	4444
3+3	7602	3690	6666

Table 8. Power supply redundancy

Configuration	Maximum power available (Watts DC)		
	2700W PSU	1350W PSU	2360W PSU
1+1	2534	1230	2222
2+1	5068	2460	4444
3+1	7602	3690	6666
4+1	10136	4920	8888
5+1	12670	6150	9922

Table 9. No Redundancy

Configuration	Maximum power available (Watts DC)		
	2700W PSU	1350W PSU	2360W PSU
1+0	2534	1230	2222
2+0	5068	2460	4444
3+0	7602	3690	6666
4+0	10136	4920	8888
5+0	11110	6150	9922
6+0	13332	7380	9922

Power supply combinations

Each power supply combination creates a customized power management behavior. This customization includes power capacity, power reserve, and power supply load strategies.

The power capacity is determined by the type of power supplies installed. A homogenous mix of 2700W power supplies provides a higher power capacity than, for example, a mix of three 2360W and three 2700W supplies.

The power reserve is a value of watts set aside in the event of a violation in the redundancy. The amount of watts set aside is a calculation of the combination and mixture of power supplies and the selected redundancy option. The power supply load strategies are adjusted to ensure the supplies are performing as efficiently as possible.

The following power supply combinations are allowed in the M1000e chassis:

- 2360W only
- 2700W only
- 1350W only
- 2360 + 2700W combination



The following power supply combinations **are not allowed** in the M1000e chassis:

- 2360W + 1350W
- 2700W + 1350W

Chassis Management Controller revisions and power supply compatibility

Table 10 lists the power supply compatibility for each CMC revision.

Table 10. Power supply compatibility per CMC revision

Power supply	CMC revision									
	2.1	2.2	2.3	3.0	4.0	4.1	4.15	4.2	4.3	4.4
2360W	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
1350W				✓	✓	✓	✓	✓	✓	✓
2700W				✓	✓	✓	✓	✓	✓	✓
DC-DC										✓

For incompatible items, the system allows the AC power to provide 12V standby power, and the AC good light displays as solid green. This does not allow the power supply to provide 12V main power and does not allow the power supply to contribute power to the chassis.

Behavior of 2700W

The CMC must be upgraded to 3.0 for 2700W power supplies to contribute power to the chassis.

Behavior of 1350W

The CMC must be upgraded to 3.0 for the 1350W to contribute power to the chassis.

Upon first chassis boot up or CMC upgrade to version 3.0 the user must acknowledge 110VAC is intended to be used. This acknowledgement can happen in the CMC GUI or through the command line interface (CLI).

Behavior of 2700W + 2360W

You can combine the 2700W and 2360W power supplies to power the M1000e enclosure, however, you must upgrade the CMC to 3.0. When mixing different capacities, the CMC adjusts the redundancy algorithm to accommodate the different capacities.

For grid redundancy, if slots 1–3 and slots 4–6 are not able to provide the same amount of power capacity, the CMC places the slots with the high capacity into standby.

For power supply redundancy, the power supply with the highest capacity is placed into redundancy.

Behavior of 1350W + 2360W/2700W

You cannot combine the 1350W (110VAC) and 2360W/2700W (220VAC) power supplies to power the M1000e chassis.

- If you insert an 110VAC power supply into a chassis running 220VAC, the system will not allow the 110VAC supply to provide capacity to the system. The AC good light will illuminate, but the DC good light will not.
- If you insert a 220VAC power supply into a chassis running 110VAC, the system will not allow the 220VAC supply to provide capacity to the system. The AC good light will illuminate, but the DC good light will not.



When starting a chassis from cold boot with a mix of 110VAC and 220VAC power supplies, the CMC will power up the first-detected power supply in slots 1-3 and slots 4-6 and will choose the lead voltage reference. All power supplies must match the lead voltage reference for the system to allow the power supplies to be powered on.

If switching from 220VAC to 110VAC or vice versa, you will need to remove all power from the chassis and power the chassis with the new voltage line.

Behavior of AC supplies with DC supplies

You cannot combine 2360W or 2700W AC power supplies with 2700W DC power supplies to power the M1000e chassis.

- If you insert a DC power supply into a chassis running on AC power, the system will not allow the DC power supply to provide capacity to the system. The AC good light will illuminate, but the DC good light will not.
- If you insert an AC power supply into a chassis running on DC power, the system will not allow the the AC power supply to provide capacity to the system. The AC good light will illuminate, but the DC good light will not.

If you wish to switch from AC power supplies to DC power supplies or vice versa, you must to remove all power from the chassis and power the chassis with the new voltage line.

Recommended number of servers

This section provides the suggested maximum number of blade servers to run under each power supply configuration.

Quarter-height blade servers

Table 11 lists the maximum number of recommended quarter-height blade servers per power supply configuration according to the following server configuration definitions:

- Light configuration: 1 x 60W processor, 4GB memory
- Medium configuration: 2 x 80W processors, 4GB memory
- Heavy configuration: 2 x 95W processors, 192GB memory

Table 11. Quarter-height servers per power supply configuration

Power supply configuration	Number of servers with 220–240V power 2700W PSUs	Number of servers with 220–240V power 2360W PSUs	Number of servers with 100–120V power 2700W/1350W PSUs ¹
2+1	22 light,	18 light,	5 light,
2+2	12 medium, or 10 heavy	10 medium, or 8 heavy	2 medium, or 2 heavy
3+0	32 light,	20 light,	13 light,
3+1	21 medium, or	18 medium, or	6 medium, or
3+3	18 heavy	15 heavy	6 heavy



Power supply configuration	Number of servers with 220–240V power 2700W PSUs	Number of servers with 220–240V power 2360W PSUs	Number of servers with 100–120V power 2700W/1350W PSUs ¹
4+1	32 light,	32 light,	21 light,
4+2	31 medium, or 26 heavy	26 medium, or 22 heavy	9 medium, or 9 heavy
5+1	32 light, 32 medium, or 32 heavy	32 light, 32 medium, or 29 heavy	29 light, 13 medium, or 13 heavy

¹With sufficient facility power available

The following information is also applicable for quarter-height blade servers:

- **Throttling** (performance reduction to conserve power):
 - 220-240V
 - > With two PSUs, throttling will occur if heavily configured blades are working very hard at the same time.
 - > With three PSUs and 95W processors, throttling will occur if heavily configured blades are a worked hard.
 - > With three PSUs, and 80W and 60W processors, throttling may occur if servers are run under a heavy load.
 - > With four or five PSUs, no throttling will occur.
 - 110V
 - > Throttling will occur in all PSU configurations if a full set of heavily configured blades are all worked hard at the same time.
- **Prioritized performance**—Under throttling conditions, there will be a performance impact on the lower-priority servers.
 - Dell recommends that you choose which servers those should be through the CMC's prioritization feature.
 - Multiple low-priority servers would distribute the impact; if all servers were same priority, then they'd each experience a slight performance impact as the throttling would be passed from server to server.
- **Delayed power-ups**—When pushing the limits of the number of servers, one to three servers might not be powered on during the initial power-up sequence if all were started simultaneously, and may be powered on in a second power-up sequence after initial blades go through their post boot power inventory and allocation.

Half-height blade servers

Table 12 lists the maximum number of recommended half-height blade servers per power supply configuration according to the following server configuration definitions:

- Light configuration: 1 x 40W processor, 2GB memory
- Medium configuration: 2 x 80W processors, 72GB memory
- Heavy configuration: 2 x 95W processors, 192GB memory



Table 12. Half-height servers per power supply configuration

Power supply configuration	Number of servers with 220–240V power 2700W PSUs	Number of servers with 220–240V power 2360W PSUs	Number of servers with 100–120V power 2700W/1350W PSUs ¹
2+1	16 light,	16 light,	9 light,
2+2	13 medium, or 10 heavy	11 medium, or 8 heavy	4 medium, or 3 heavy
3+0		16 light,	16 light,
3+1		16 medium, or	8 medium, or
3+3		15 heavy	6 heavy
4+1	Any server combination		16 light,
4+2			13 medium, or 10 heavy
5+1		Any server combination	16 light, 16 medium, or 13 heavy

¹With sufficient facility power available

The following information is also applicable for half-height blade servers:

- **Throttling** (performance reduction to conserve power):
 - 220–240V power:
 - > 2 PSUs — Throttling will occur if heavily configured blades are worked very hard at the same time.
 - > 3 PSUs:
 - 130W processors — Some throttling will occur if heavily configured blades are worked hard at the same time.
 - 95W or 80W processors — Throttling may happen if all servers are run under a very heavy load.
 - > 4 or 5 PSUs — No throttling will occur
 - 110V power — Throttling will occur in all power supply configurations if a full set of heavily configured blades are worked hard at the same time.
- **Prioritized performance** — Under throttling conditions, there will be a performance impact on the lower-priority servers.
 - Dell recommends you choose which servers are prioritized through CMC’s prioritization feature.
 - Multiple low-priority servers can spread out the impact of prioritization. If all servers are the same priority, they would each experience a slight performance impact as the throttling would be passed from server to server.
- **Delayed power-ups** — When pushing the limits of the guidelines listed in Table 12, one to three servers may not be powered on during the initial power-up sequence if all servers were started simultaneously. The servers may need to be powered on in a second power-up action after the initial blades go through their post-boot power inventory and allocation.



Dell recommends that all customers running a power-hungry configuration should install the modular iDRAC with Lifecycle Controller firmware 3.02 upgrade, which contains the following enhancements:

- Improved calculation for power budgeting to enable more servers to be powered up simultaneously.
- Functionality to enable iDRAC with Lifecycle Controller to retry power up for blade servers automatically in the event that some servers do not power up initially due to the pre-boot power inventory.

Full-height blade servers

Table 13 lists the maximum number of recommended half-height blade servers per power supply configuration according to the following server definitions:

- Light configuration: 2 x 95W processor, 4 DIMMs
- Medium configuration: 4 x 105W or 2 x 130W processors, 16 DIMMs
- Heavy configuration: 4 x 105W or 2 x 130W processors, 32 DIMMs

Table 13. Full-height servers per power supply configuration

Power supply configuration	Number of servers with 220–240V power 2700W PSUs	Number of servers with 220–240V power 2360W PSUs	Number of servers with 100–120V power 2700W/1350W PSUs ¹
2+1	8 light,	8 light,	3 light,
2+2	7 medium, or 5 heavy	5 medium, or 4 heavy	2 medium, or 1 heavy
3+0		8 light,	6 light,
3+1		8 medium, or	4 medium, or
3+3		7 heavy	3 heavy
4+1	Any server combination		8 light,
4+2			6 medium, or 4 heavy
5+1		Any server combination	8 light, 8 medium, or 6 heavy

¹With sufficient facility power available

The following information is applicable for full-height blade servers:

- **Throttling** (performance reduction to conserve power):
 - 220-240V power:
 - > 2 PSUs — Throttling will occur if heavily configured blades are worked very hard at the same time.
 - > 3 PSUs:
 - 4 processors — Some throttling will occur if heavily configured blades are worked hard at the same time. Minimal throttling is expected for 16-DIMM configurations.
 - 2 processors — Throttling may occur if all servers are run under a very heavy load and have 130W processors and/or 32 DIMMs populated. Minimal throttling is expected with 95W processors and/or 16 DIMMs.



- > 4 or 5 PSUs — No throttling will occur.
 - 110V power — Throttling will occur in all power supply configurations if a full set of heavily configured blades are worked hard at the same time.
- **Prioritized performance** — Under throttling conditions, there will be a performance impact on the lower-priority servers.
 - Dell recommends using the CMC Power Prioritization feature to indicate which servers have the highest priority and which have the lowest priority.
 - Multiple low-priority servers can spread out the impact of prioritization. If all servers are the same priority, they would each experience a slight performance impact as the throttling would be passed from server to server.
- **Delayed power-ups** — When pushing the limits of the guidelines listed in Table 13, one to three servers may not be powered on during the initial power-up sequence if all servers were started simultaneously. The servers may need to be powered on in a second power-up action after the initial blades go through their post-boot power inventory and allocation. iDRAC with Lifecycle Controller 3.02 enables automatic retries if blades attempt simultaneous power-ups and do not power on.

For the PowerEdge M910 blade server, Dell recommends that you install the modular iDRAC with Lifecycle Controller firmware 3.02 upgrade, which contains the following enhancements:

Improved power budgeting to lower the amount of power budget that the M910 requests from the CMC, enabling more servers to run on fewer power supplies.

Functionality to enable iDRAC with Lifecycle Controller to retry power-ups for blade servers automatically in the event that some servers do not power up initially due to the pre-boot power inventory.

Power management

Power is no longer just about power delivery, it is also about power management. The M1000e system offers many advanced power management features. Most of these features operate transparently, while others require only a one-time selection of desired operating modes.

Shared power takes advantage of the large number of resources in the modular server, distributing power across the system without the excess margin required in dedicated rack-mount servers and switches. The M1000e has an advanced power budgeting feature, controlled by the CMC and negotiated in conjunction with the iDRAC with Lifecycle Controller on every server module. Prior to any server module powering up, through any of its power-up mechanisms such as AC recovery, WOL or a simple power button press, the server module iDRAC with Lifecycle Controller performs a sophisticated power budget inventory for the server module, based upon its configuration of processors, memory, I/O and local storage. Once this number is generated, the iDRAC with Lifecycle Controller communicates the power budget inventory to the CMC, which confirms the availability of power from the system level, based upon a total chassis power inventory, including power supplies, iKVM, I/O modules, fans, and server modules. Since the CMC controls when every modular system element powers on, it can set power policies on a system level.

In coordination with the CMC, iDRAC with Lifecycle Controller hardware constantly monitors actual power consumption at each server module. This power measurement is used locally by the server module to ensure that its instantaneous power consumption never exceeds the budgeted amount. While the system administrator may never notice these features in action, what they enable is a more aggressive utilization of the shared system power resources. Thus the system is never flying blind in regard to power consumption, and there is no danger of exceeding power capacity availability, which could result in a spontaneous activation of power supply over-current protection without these features.



The system administrator can also set priorities for each server module. The priority works in conjunction with the CMC power budgeting and iDRAC power monitoring to ensure that the lowest priority blades are the first to enter any power optimization mode, should conditions warrant the activation of this feature.

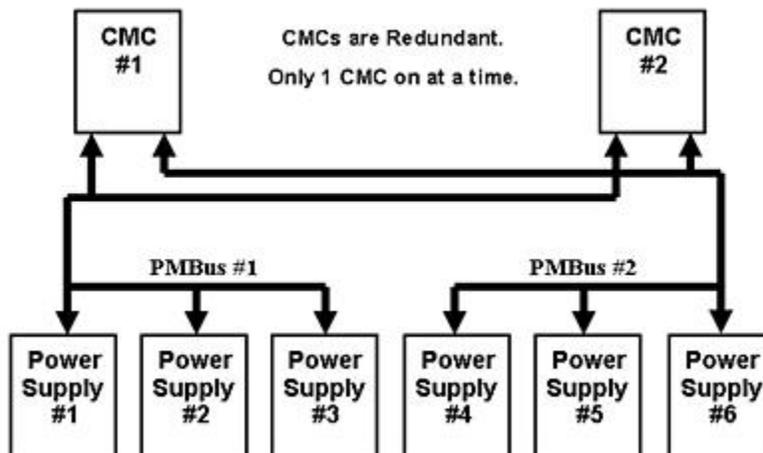
Power capping is set at the chassis level for our blade servers and not at the blade-server level, so components like a processor or memory can throttle down when necessary on lower priority blade servers. An allocation is taken out for a component (fans, I/O modules), the remainder is applied to the blades, and then throttling is applied if required to get under the cap. If all the blades are set up with the same priority, they start throttling down the processor, memory, and so on. A variety of BIOS settings throttle the processor or not, depending on load.

If power consumption demands exceed available power, the enclosure throttles back the power supplied to blades as prioritized in the CMC. The blades do not shut down; rather they slow down if necessary; Dell designed the system this way on purpose, in response to customer feedback that they did not want the blades to shut themselves down under any condition. I/O modules, on the other hand, are shut down prior to permanent damage, as they are less tolerant to power variation than the blade server hardware.

The M1000e is compliant with the PMBus Specification 1.1 using this power management standard for status, measurement, and control. The M1000e power supplies continuously monitor AC input current, voltage and power, enabling exposure of data to enterprise-level management tools, such as Dell's OpenManage Power Center. Real-time power consumption is viewable per system.

For more information on Dell OpenManage Power Center, visit [Dell OpenManage Power Center](#).

Figure 22. PMBus communication channels



Each power supply communicates information about itself through the PMBus (I2C bus):

- Vendor part number
- Revision serial number
- Power capability

Each power supply performs the following functions:

- Control for the main power regulator
- Reports power-supply status across the PMBus
 - Input/Output Voltage/Current
 - AC power consumption



- DC power consumption
- Error status
- Individual Status Signals from the power supply to the CMC
- AC Good
- DC Good
- PSU Present Alert

VMware has included consuming the current power consumption and current power cap/limit retrieval through Dell-specific IPMI commands through iDRAC. They are using this to report the total power consumed by the server and also using this as part of their calculations to determine/approximate the VM-level power. Unfortunately, there is no integration as yet of generating the view/consolidation of power attributes across all the blades in a chassis, such as a chassis view; or the ability make use of that knowledge (available through SMBIOS structures) to make use of that in DPM/DRS or other tools/applications.

For more information, see the *Dell Chassis Management Controller Firmware User's Guide* on Dell.com/Support/Manuals.

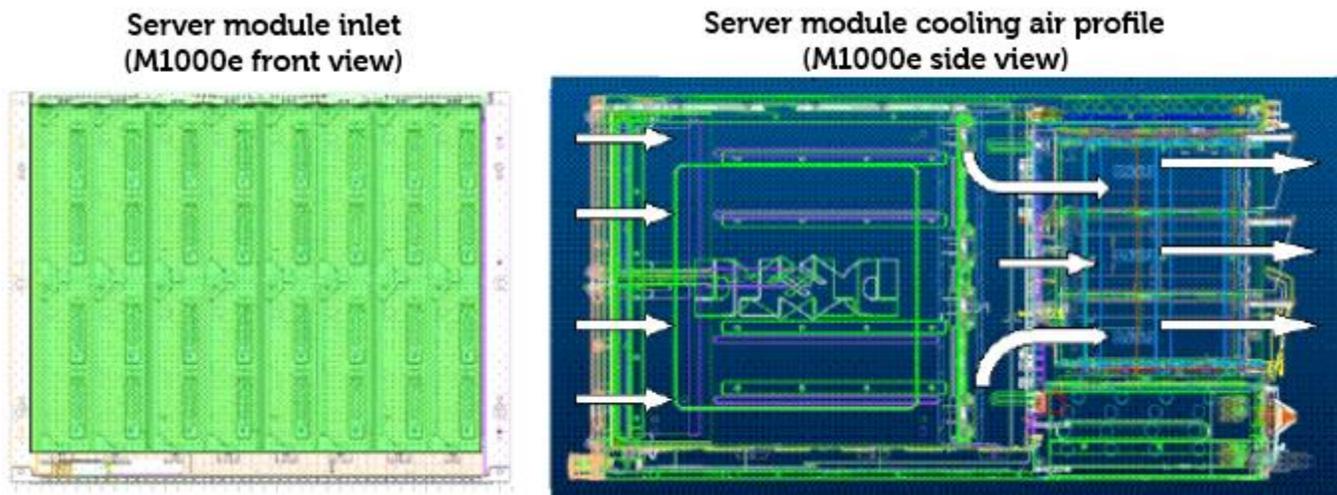
Heat dissipation

The cooling strategy for the M1000e supports a low-impedance, high-efficiency design philosophy. Driving lower airflow impedance allows the M1000e to draw air through the system at a lower operating pressure and reduces the system fan power consumed to meet the airflow requirements of the system.

The low impedance design is coupled with a high-efficiency, air-moving device designed explicitly for the PowerEdge M1000e chassis. The efficiency of an air-moving device is defined as the work output of the fan as compared to the electrical power required to run the fan. The M1000e fan operates at extreme efficiencies, which correlates directly into savings.

The high-efficiency design philosophy also extends into the layout of the subsystems within the M1000e. The server modules, I/O modules, and power supplies are incorporated into the system with independent airflow paths. This isolates these components from pre-heated air, reducing the required airflow consumptions of each module.

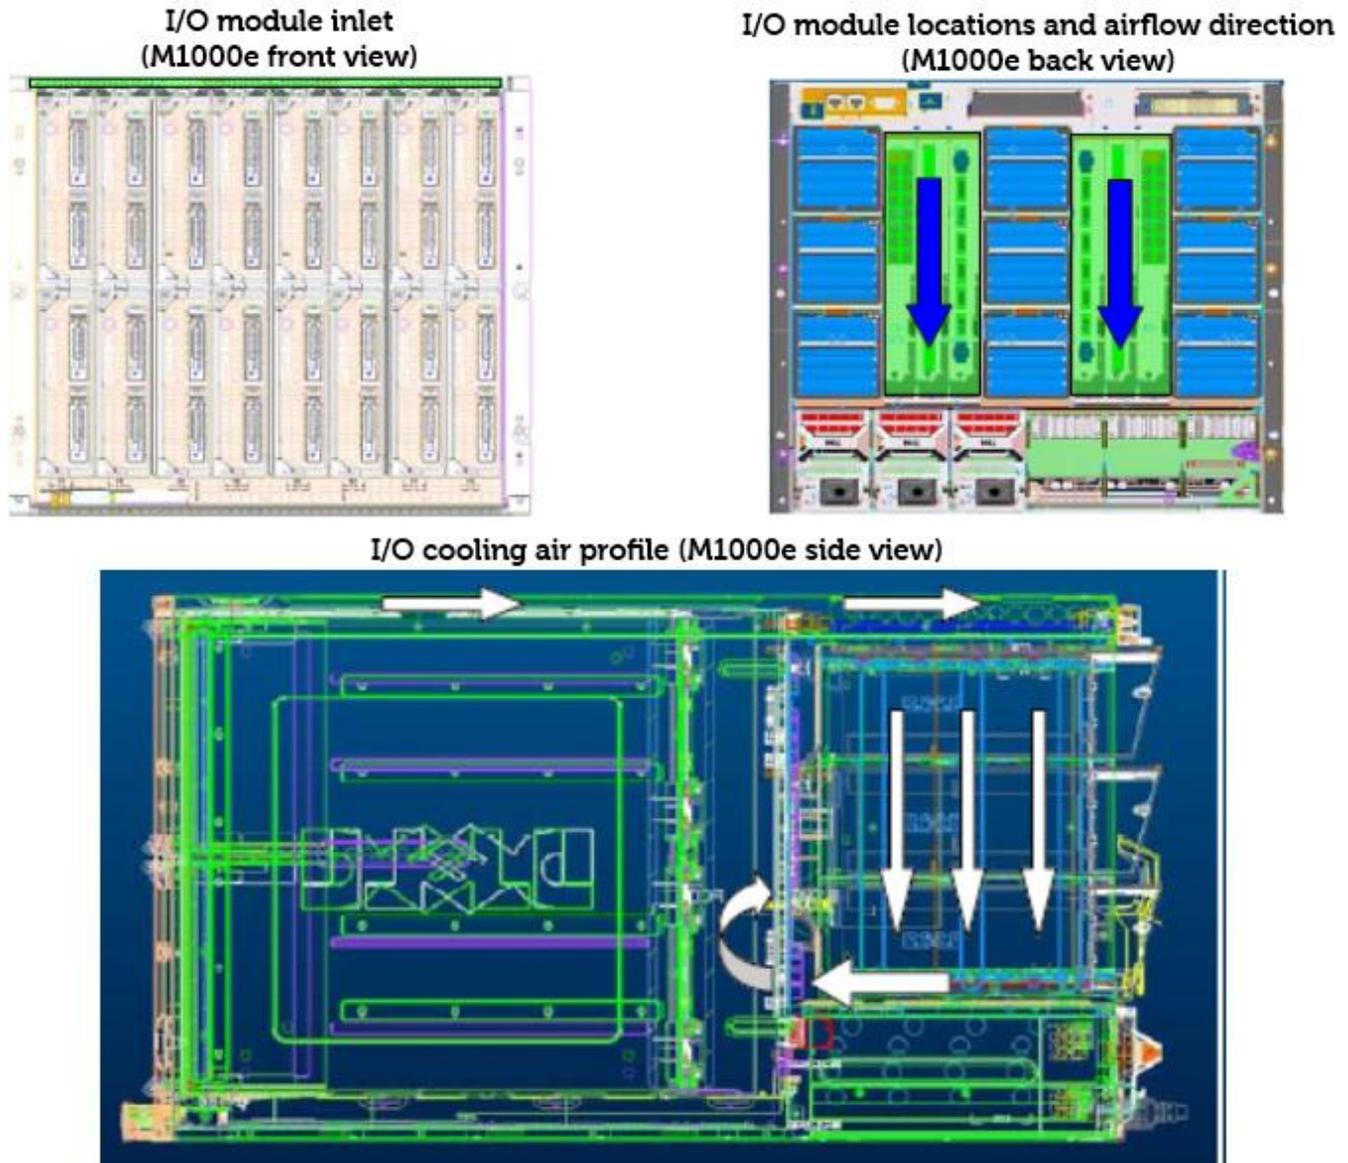
Figure 23. Server cooling air profile



The server modules are cooled with traditional front-to-back cooling. As shown in Figure 23, the front of the system is dominated by inlet area for the individual server modules. The air passes

through the server modules, through venting holes in the midplane, and is then drawn into the fans that exhaust the air from the chassis. There are plenums both upstream of the midplane, between the midplane and the blades, and downstream of the midplane, between the midplane and the fans, to more evenly distribute the cooling potential from the three columns of fans across the server modules.

Figure 24. I/O module cooling air profile

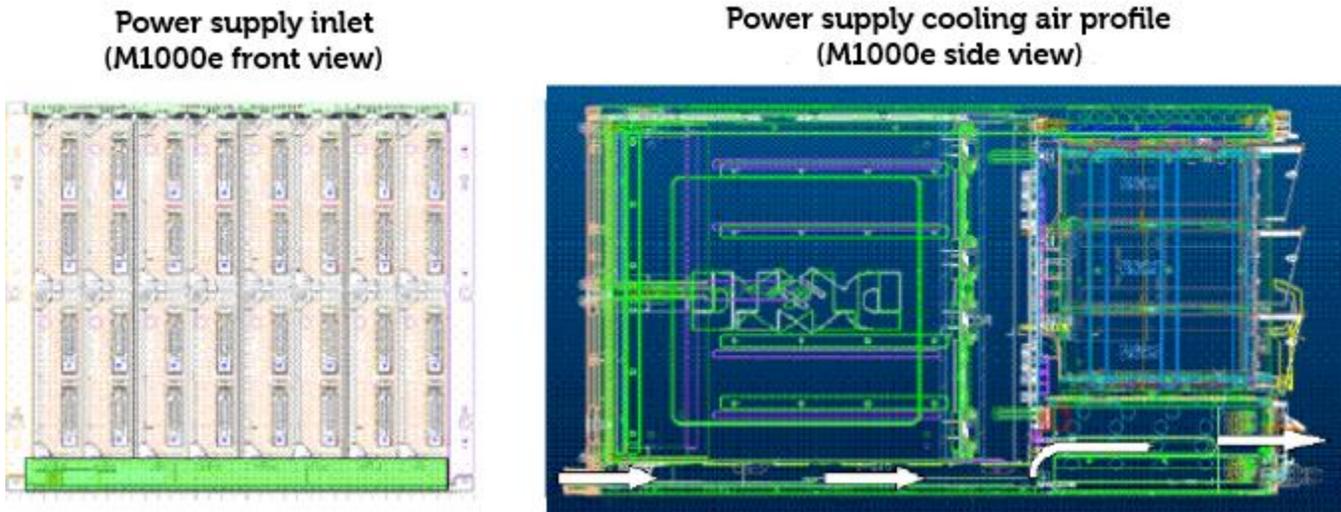


The I/O modules use a bypass duct to draw ambient air from the front of the system to the I/O module inlet, as seen in Figure 24. This duct is located above the server modules. This cool air is then drawn down through the I/O modules in a top to bottom flow path and into the plenum between the midplane and fans, from where it is exhausted from the system.

The power supplies, located in the back of the system, use basic front-to-back cooling, but draw their inlet air from a duct located beneath the server modules, as seen in Figure 25. This ensures that the power supplies receive ambient temperature air.



Figure 25. Power supply inlet and cooling air profile



This hardware design is coupled with a thermal cooling algorithm that incorporates:

- Server module level thermal monitoring by the iDRAC with Lifecycle Controller
- I/O module thermal health monitors
- Fan control and monitoring by the CMC

The iDRAC on each server module calculates the amount of airflow required on an individual server module level and sends a request to the CMC. This request is based on temperature conditions on the server module, as well as passive requirements due to hardware configuration. Concurrently, each IOM can send a request to the CMC to increase or decrease cooling to the I/O subsystem. The CMC interprets these requests, and can control the fans as required to maintain server and I/O module airflow at optimal levels.

Acoustics

The M1000e is engineered for sound quality in accordance with the Dell Enterprise acoustical specification. Compared to previous generations of products, the fans have more levels of control and finer tuning of the fan behavior. Firmware is optimized to choose the lowest fan speeds and therefore the lowest acoustical output for any configuration (components installed), operating condition (applications being run), and ambient temperature. Because acoustical output is dependent and indeed minimized for each combination of these variables, no single acoustical level (sound pressure level or sound power level) represents the M1000e, and instead boundaries on sound power level are provided below:

- Lowest Fan Speed: Upper Limit A-weighted Sound Power Level, LwA-UL, is 7.5 bels.
- Full Fan Speed: Upper Limit A-weighted Sound Power Level, LwA-UL, is 9.7 bels.
- LwA-UL is the upper limit sound power level (LwA) calculated per section 4.4.2 of ISO 9296 (1988) and measured in accordance to ISO 7779 (1999).
- Acoustical models have been provided to predict performance between these bounds in the Dell Energy Smart Solution Advisor (ESSA) tool. Learn more at Dell.com/ESSA.

A few things to be aware of:

- Fans are loud when running at full speed. It is rare that fans need to run at full speed. Ensure that components are operating properly if fans remain at full speed.

- The CMC automatically raises and lowers the fan speed to a setting that is appropriate to keep all modules cool.
- If a single fan is removed, all fans are set to 50% speed, if the enclosure is in Standby mode. If the enclosure is powered on, removal of a single fan is treated like a failure (nothing happens).
- Re-installation of a fan causes the rest of the fans to settle back to a quieter state.
- Whenever communication to the CMC or iDRAC is lost, for example during a firmware update, the fan speed is increased and creates more noise.



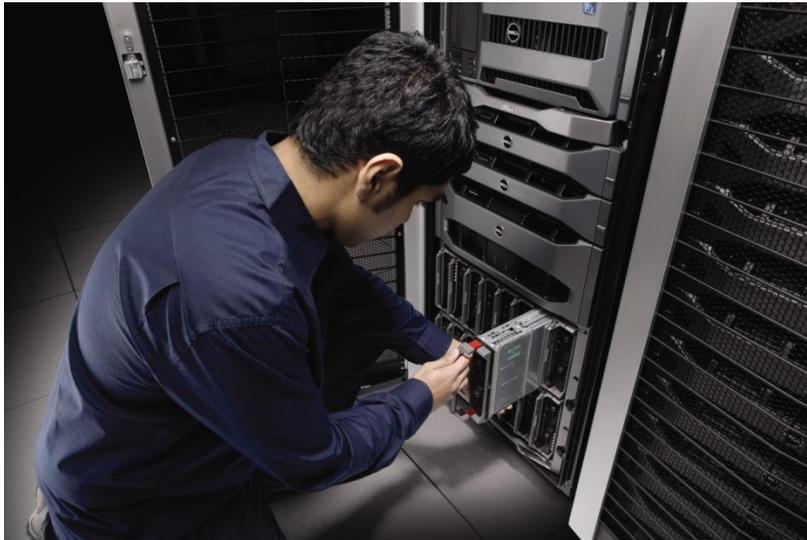
9 Rack information

Rack enclosure features for the PowerEdge M1000e include:

- Complete zero-U (0U) power management options
- Split back doors for easy access
- Complete markings on all posts for easy installation
- Full (42U) and half-rack (24U) configurations available
- Tool-less rail installation
- Enhanced tool-less cable management
- Complete data center configuration tool for power and environmental calculations

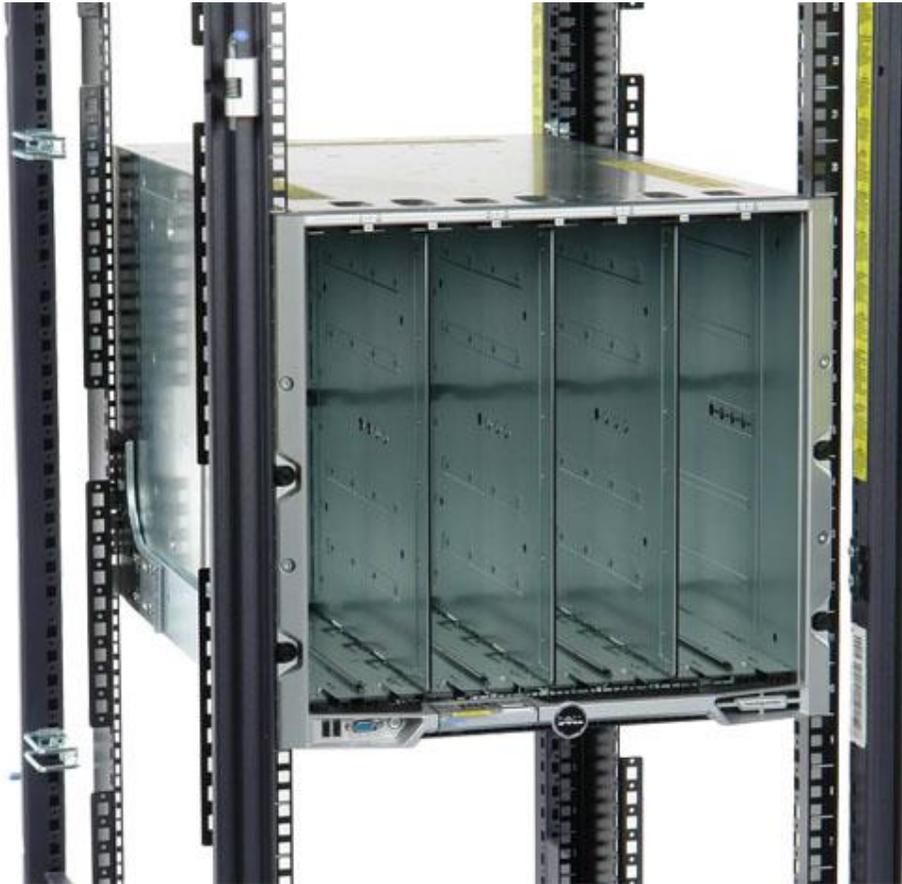
The M1000e is installed on the rails by simply resting the back of the system on the rail ledges, pushing the system forward until it fully seats, and tightening the thumbscrews on the chassis front panel.

Figure 26. M1000e mounted in a rack



When properly installed in a rack, the enclosure sits on a sturdy pair of rails and is held in place by four thumbscrews.

Figure 27. Empty enclosure in a rack



The RapidRails static rail system for the M1000e provides tool-less support for racks with square mounting holes including all generations of Dell racks except for the 4200 and 2400 models. Also available are the VersaRails static rails, which offer tooled mounting support for racks with square or unthreaded round mounting holes. Both versions include a strain relief bar and cable enumerator clips to help manage and secure the cables exiting the back of the system.

Rails

The RapidRails and VersaRails static rail kits for the M1000e are identical, except for their mounting bracket designs. The mounting brackets on the RapidRails static rails have hooks and a lock button for supporting tool-less installation in 4-post racks with square mounting holes.

RapidRails static rails for square-hole racks support the following:

- Tool-less installation in 19-inch EIA-310-E compliant square-hole 4-post racks, including all generations of Dell racks except for the 4200 and 2400 series
- Strain relief bar and cable enumerators for managing and securing cables
- Square-hole rack adjustment range of 712 to 755mm
- Minimum rail depth of 703mm



Figure 28. RapidRails static rails

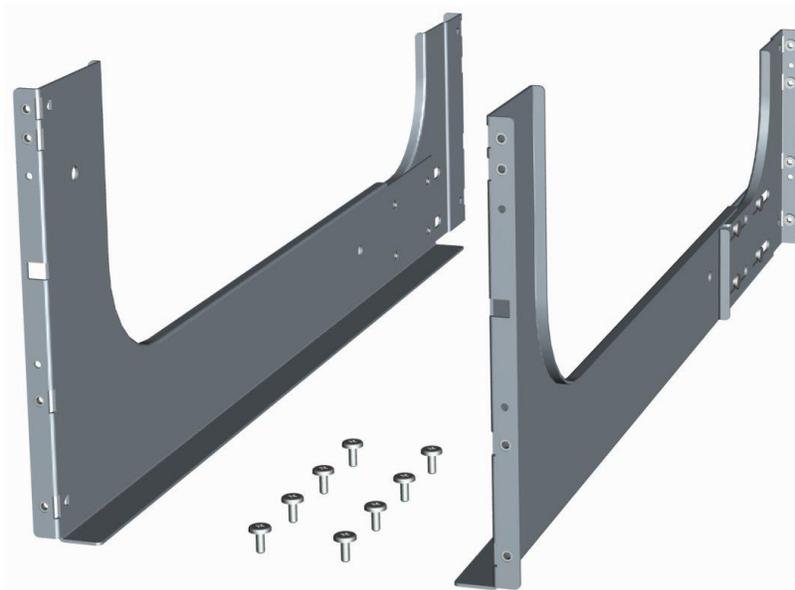


VersaRails static rails for square-hole or round-hole racks support the following:

- Tooled installation in 19-inch EIA-310-E compliant square or unthreaded round-hole 4-post racks
- Strain relief bar and cable enumerators for managing and securing cables
- Square-hole rack adjustment range of 706 to 755mm
- Round-hole rack adjustment range of 706 to 755mm
- Minimum rail depth of 703mm

The mounting brackets on the VersaRails static rails have threaded clinch nuts rather than hooks and a lock button in order to support tooled installation in 4-post racks with unthreaded round mounting holes. The VersaRails static rails can also be mounted in square-hole racks if desired.

Figure 29. VersaRails static rails



The VersaRails static rails are not intended to be mounted in threaded-hole racks, because the rails cannot be fully tightened and secured against the rack mounting flange. Neither the VersaRails nor the RapidRails kits support mounting in 2-post racks. The min-max values listed in Table 14 represent the allowable distance between the front and back mounting flanges in the rack.

Table 14. Supported racks

Mounting interface	Rail type	Rack types supported					Rail adjustability range (mm)						Rail depth (mm)	
		4-Post			2-Post		Square		Round		Threaded		w/out CMA	with CMA
		Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max		
RapidRails	Static	✓	✗	✗	✗	✗	712	755	—	—	—	—	703	—
VersaRails	Static	✓	✓	✗	✗	✗	706	755	706	755	—	—	703	—

Cabling

One of the advantages of a modular server system is the reduction in cable management needs within a rack system. The inclusion of fabric switches, integrated KVM, and system management aggregation at the CMCs provides six-fold or better cable reduction. Table 15 shows a comparison of a typical reduction available when using the M1000e modular system with integrated switches, compared to traditional rack-and-stack components. The configuration in Table 15 assumes a server with four Ethernet ports and two Fibre Channel ports. In support of the M1000e, Dell offers a modular system cable management system to ease installation in Dell or other industry-standard racks.

Table 15. Typical modular server system rack height and cable reduction

Component	Rack height	AC power cables	Ethernet cables	FC cables	KVM cables
2-socket server	1Ux16	2x16	4x16	2x16	USBx16 + VGAx16
KVM	1U	1	—	—	USBx1 + VGAx1
Ethernet switches	1Ux4	1x4	—	—	—
FC switches	1Ux2	1x2	—	2x2	—
Total rack	23U height	39 AC cables	72 Ethernet cables	36 FC cables	USBx17 + VGAx17
M1000e equivalent	10U height	6 AC cables	16 Ethernet cables	4 FC cables	USBx1 + VGAx1

Cable management arm (CMA)

Since the M1000e does not need to be extended out of the rack for service, neither the RapidRails nor the VersaRails static rail systems support a cable management arm. Included instead are a strain relief bar and cable enumerator clips (Figure 30) to help manage and secure the potentially large number of cables exiting the back of the system.

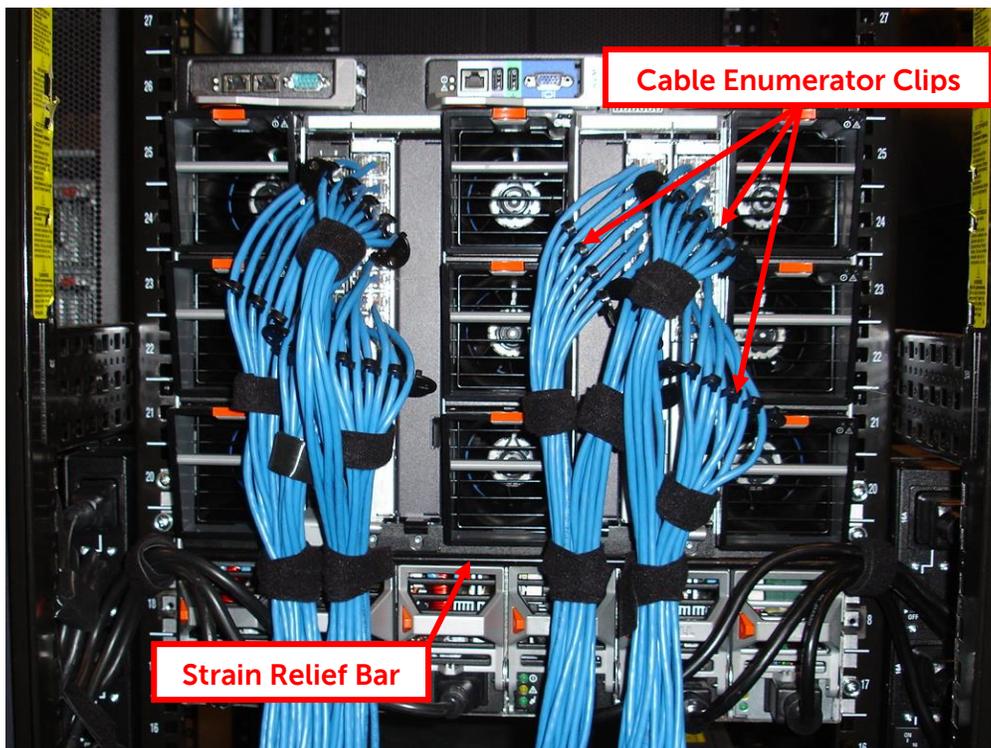


Figure 30. M1000e strain relief bar and cable enumerator clip (12 per kit)



The strain relief bar and cable enumerator clips can be used to help manage and secure the power cords and I/O cables exiting the back of the system as shown in Figure 31.

Figure 31. M1000e strain relief bar and cable enumerator clips



For information on cabling for the M1000e modular system, see the *Rack Installation Guide* on Dell.com/Support/Manuals.



Rack and Stack

The PowerEdge M1000e has been qualified for Rack and Stack support along with specified blade servers. Details are as follows:

- Qualifications are limited to Dell branded racks shipped by an air-ride van.
- The M1000e is qualified to ship fully populated with full height blades up to M910 blade weight or up to pallet weight limits.
 - The heaviest blade qualified to date is the M910.
 - Any future blade that is heavier than the M910 or incorporates significant changes in mechanical layout may require qualification testing.
- The M1000e is qualified to ship partial configurations of blades up to pallet weight limits.
- The M1000e is qualified to ship without blades.
- The M1000e is qualified to ship on VersaRail only.
- The following screws should be secured with permanent Loctite™ and seated to the following torques:
 - VersaRail screws 20 in-lb
 - Chassis thumbscrews 20 in-lb
 - Screws that secure the shipping clip 20 in-lb
 - Screws that secure the strain relief bar 20 in-lb
- I/O modules should have their latches secured mechanically.
- Power supply blanks should be shipped uninstalled.
- The M1000e PDUs are qualified for zero-U tool-less mounting.



10 Virtualization

The M1000e and the blade servers it contains have been designed for optimal use with all major virtualization software platforms.

Figure 32. Examples of major virtualization platforms supported



- **VMware® vCenter™ 4.0/ESXi v4.0**
- **x86 Virtualization Software Market leader**
- **Enterprise-class feature set**



- **Familiar user experience backed by Microsoft**
- **Comprehensive management for physical and virtual environments**
- **Cost-effective solution for Microsoft environments**



- **Citrix® XenServer™ Express & Enterprise**
- **Cost-effective and easy to use, based on open source Xen®**
- **Rich feature set**



- **Software solutions to facilitate virtualization with Dell hardware**
- **PowerRecon for data center performance optimization**
- **PowerConvert® to assist in efficiently managing movement of server workloads between physical and virtual platforms**



- **Processing Area Network (PAN) Manager**
- **Infrastructure virtualization software enables customers to create and manage a single pool of physical and virtualized servers**

The M1000e platform offers many benefits for virtualization:

- Data center consolidation
 - High density form factor
 - I/O bandwidth and switch port savings
 - Large memory capacity
- Ease-of-management and deployment
 - Management options
 - I/O virtualization
 - Chassis LCD display
 - Embedded hypervisor
- Reduced downtime
 - Persistent addresses
 - Fully redundant power and cooling
 - Fully redundant I/O
 - Hot-pluggable drives
- Power and cooling efficiency
 - Super-efficient power supplies
 - Optimized airflow
 - Best-in-class fan technology



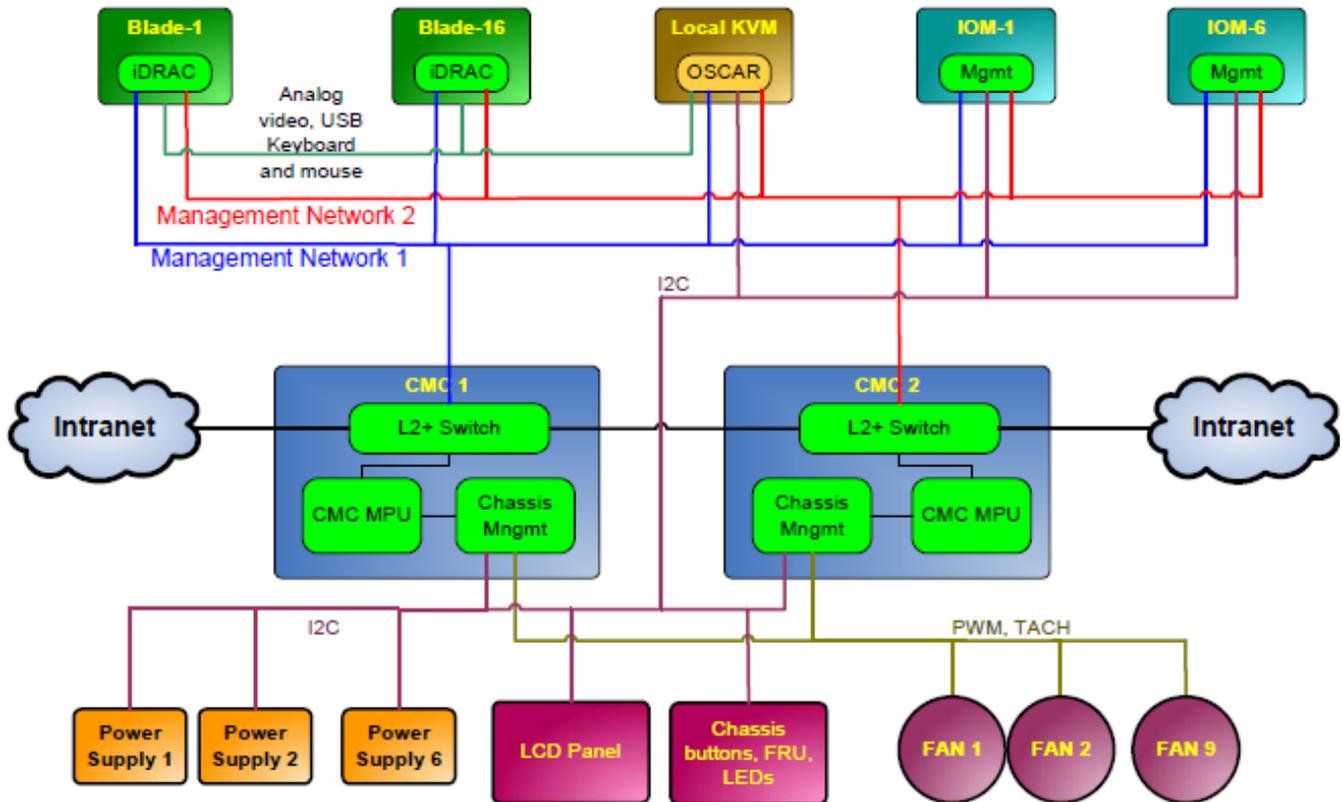
11 Systems management

The M1000e server solution offers a holistic management solution designed to fit into any data center. The solution offers the following features:

- Dual-redundant chassis management controllers (CMC) when redundant CMC option is chosen
 - Powerful management for the entire enclosure
 - Includes: real-time power management and monitoring, flexible security, status/inventory/alerting for blades, I/O and chassis
- iDRAC
 - One per blade with full DRAC functionality, similar to other Dell servers including vMedia/KVM
 - Integrates into CMC or can be used separately
- iKVM
 - Embedded in the chassis for easy KVM infrastructure incorporation, allowing one administrator per blade
 - Control panel on front of M1000e for crash-cart access
- Front LCD
 - Designed for deployment and local status reporting

The onboard graphics and keyboard/mouse USB connect to an optional system-level iKVM module for local KVM access. Full USB access is available through the server module front panel.

Figure 33. System Management Architecture Diagram



Management connections transfer health and control traffic throughout the chassis. The system management fabric is architected for 100BaseT Ethernet over differential pairs routed to each module. There are two 100BaseT interfaces between CMCs, one switched and one unswitched. All system management Ethernet is routed for 100 Mbps signaling. Every module has a management network link to each CMC, with redundancy provided at the module level. Failure of any individual link will cause failover to the redundant CMC.

Server management

See the *Dell PowerEdge Mxxx Technical Guide* on Dell.com/PowerEdge for each of the compatible blade servers offered for more details on server management.

Enclosure management

The CMC provides secure remote management access to the chassis and installed modules. The M1000e must have at least one CMC and supports an optional redundant module, each occupying a slot accessible through the back of the chassis. Redundancy is provided in an Active — Standby pairing of the modules and failover occurs when the active module has failed or degraded. The CMC interfaces through dual stacking 10/100/1000 Ethernet ports and one serial port. The CMC serial port interface provides common management of up to six I/O modules through a single connection.

The CMC provides many features, including:

- Deployment
 - LCD-based deployment wizard
 - Single secure interface for inventory, configuration, monitoring, and alerting for server modules, chassis infrastructure and I/O Modules
 - Centralized configuration for iDRAC, I/O Modules and CMC
 - One-to-many iDRAC configuration
 - One-to-many blade boot device selection
 - One-to-many vMedia file share
 - Customized slot naming
 - I/O module configuration and launch
 - WWN/MAC display and persistence with FlexAddress; manages FlexAddress ports
 - Support for Network Time Protocol (NTP)
- Monitoring and troubleshooting
 - User interface entry point (web, telnet, SSH, serial)
 - Monitoring and alerting for chassis environmental conditions or component health thresholds. This includes but is not limited to the following:
 - > Real-time power consumption
 - > Power supplies
 - > Fans
 - > Power allocation
 - > Temperature
 - > CMC redundancy
 - I/O fabric consistency
 - Consolidated status reporting and event logs
 - > Email and SNMP alerting
 - > Support for remote syslog
 - > Blade events displayed in CMC
 - Consolidated chassis, blade, and I/O Inventory



- Virtual Server Reseat simulates blade removal and insertion
- Remotely blink LEDs to identify components

Figure 34. Web interface for the chassis management controller



- Updating
 - Reporting of firmware versions
 - One-to-many iDRAC firmware update
 - Consolidated CMC and iKVM firmware update
 - One-to-many update of drivers and firmware through a remote file share (with Repository Manager)
 - 2 x 10/100/1000Mb Ethernet ports + 1 serial port
- Real-time power and thermal monitoring and management
 - Consolidated chassis and blade power reporting
 - Power budget management and allocation (requires iDRAC7 Enterprise)
 - Real-time system AC power consumption with reset-able peak and minimum values
 - System level power limiting and slot based power prioritization
 - Management for dynamic power engagement functionality
 - Management fan speed control
 - Power sequencing of modules in conjunction with the defined chassis power states
- Separate management network
 - Configuration of the embedded management switch, facilitating external access to manageable modules
 - Connection from management network to iDRAC on each of the blades and the management interfaces on the integrated I/O Modules
 - Second Ethernet port, supporting daisy chaining of CMCs for improved cable management
- Security
 - Local authentication and AD/LDAP integration



- Support for multiple levels of user roles and permissions for control of chassis, I/O, and server blades, including Microsoft® Active Directory®
- IPv6 support
- VLAN tagging for iDRAC, CMC, and IOMs
- Two factor authentication with SmartCard
- Single sign on using OS credentials (with AD)
- Private Key Authentication (PK Auth)
- Secure Web (SSL) and CLI (Telnet/SSH) interfaces
- Support for industry-standard interfaces
 - SMASH CLP
 - WSMAN
 - CIM XML
 - SNMP

The iDRAC on each server module is connected to the CMC through dedicated, fully redundant 100 Mbps Ethernet connections wired through the midplane to a dedicated 24-port Ethernet switch on the CMC, and exposed to the outside world through the CMC's external Management Ethernet interface (10/100/1000M). This connection is distinct from the three redundant data Fabrics A, B and C. Unlike previous generations of Dell server modules, the iDRAC's connectivity is independent of, and in addition to, the onboard GbE LOMs on the server module. Each server module's iDRAC has its own IP address and can be accessed, if security settings allow, directly through a supported browser, telnet, SSH, or IPMI client on the management station.

Figure 35. Chassis management controller



The CMC 3.0 release added:

- An improved user interface
- LDAP support
- Auto propagation of OS hostname to slot name

The CMC 3.1 release added:

- Clearly articulated error messages
- Enhanced power logging to Syslog server
- CMC configuration backup and restore
- New power option to favor performance over redundancy
- OMSA quick-launch link
- Multi-chassis Visibility



The CMC 3.2 release added:

- Firmware update of server component(s) across multiple servers in the chassis using Lifecycle Controller.
- Extended Non-Volatile (persistent) Storage, using 2 GB Secure Digital (SD) media in the CMC card slot(s).
- Power Cap can be set to 16685W AC.
- Chassis Group Quick Launch and Single-Sign On (SSO) from the leader to iDRAC GUI and KVM.

The CMC 4.0 added:

- Blade server BIOS capture and one-to-many apply
- Cloning of the lead chassis configuration to members
- Cluster-wide inventory report makes it easy to find empty slots or blades with a specific configurations of CPU and memory

The CMC 4.2 added:

- Enables initial PowerEdge M I/O Aggregator deployment for server administrators in CMC and from consoles

The CMC 4.3 added:

- Extended one-to-many firmware update to include update capability for supported 12th generation Fiber Channel (FC) mezzanine cards
- Enables update of Dell PowerEdge M IO Aggregator firmware from the CMC GUI
- Enables saving the blade server BIOS capture to hard disk for restoring to the same or a different chassis
- RACADM support to configure Dell EqualLogic PS-M4110 Blade Array
- Multi-chassis management
 - Capability to select chassis configuration properties from Lead Chassis and push to Group Members
 - Capability for Group Members to keep their chassis settings synchronized with the Lead Chassis
- Capability to reset iDRAC7 without rebooting the blade server from the CMC
- Default key size for Certificate Signing Requests changed from 1024 to 2048

The CMC 4.4 added:

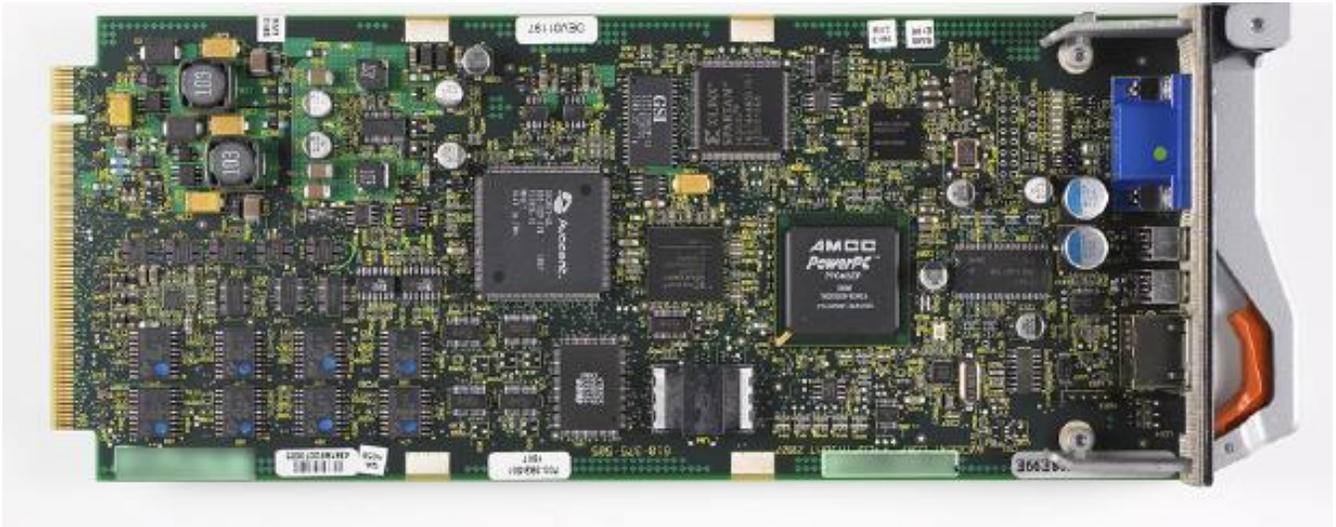
- Support for DC power supplies
- Blade server configuration capture
 - Expanded to include multiple server components such as iDRAC7 with Lifecycle Controller, and NIC
 - Can be assigned to a chassis slot and applied on server insertion
- Enhanced reporting for Dell PowerEdge M I/O Aggregator to include FCoE, uplink and downlink status, and stacking information

Dell recommends upgrading to the latest version of CMC to ensure optimal performance over the widest possible range of blade server configurations. For more information on CMC features, see the *Dell PowerEdge M1000e Enclosure Owner's Manual* on Dell.com/Support/Manuals.



Integrated keyboard and mouse controller

Figure 36. M1000e iKVM



The M1000e modular enclosure supports one optional iKVM module. This module occupies a single slot accessible through the back of the chassis. The iKVM redirects local server module video, keyboard, and mouse electrical interfaces to either the iKVM local ports or the M1000e front panel ports. The iKVM allows connection to a VGA monitor, USB keyboard, and USB mouse without use of a dongle. The iKVM also has an Analog Console Interface (ACI) compatible RJ45 port that allows the iKVM to tie the interface to a KVM appliance upstream of the iKVM through CAT5 cabling. Designed with Avocent technology, the ACI port reduces cost and complexity by giving access for sixteen servers using only one port on an external KVM Switch.

The iKVM contains a seventeenth-blade feature, connecting to the CMC Command Line Interface through the KVM switch and allowing text-based deployment wizards on VGA monitors. iKVM firmware is updated through the CMC.

For more information on iKVM module features, see the *Dell PowerEdge M1000e Enclosure Owner's Manual* on Dell.com/Support/Manuals.

The front of the enclosure includes two USB connections for a keyboard and mouse, along with a video connection port, both of which require the Avocent iKVM switch to be activated for them to be enabled. These ports are designed for connecting a local front crash-cart console to be connected to access the blade servers while standing in front of the enclosure.

Figure 37. Front keyboard/video ports



Dell modular servers also include vKVM as a standard feature, routing the operator's keyboard output, mouse output and video between the target server module and a console located on the system management IP network. With up to two simultaneous vKVM sessions per blade, remote management now satisfies virtually any usage model. vMedia is also now standard, providing emulation of USB DVD R/W, USB CD R/W, USB flash drive, USB ISO image, and USB floppy over an IP interface. Connection to vKVM and vMedia is through the CMC, with encryption available on a per stream basis.

It is possible to connect the following Dell and Avocent KVMIP switches to the iKVM card in the M1000e blade enclosure using a CAT5 cable.

- Dell:
 - 2161DS-2
 - 4161DS
 - 2321DS
 - 180AS
 - 2160AS
- Avocent:
 - All DSR xx20, xx30, xx35 models
 - All Mergepoint Unity® models

For other Avocent branded models, connect to the card using the USB adapter.

For more information on the iKVM see the *Dell PowerEdge M1000e Enclosure Owner's Manual* on and the *Dell Chassis Management Controller Firmware User's Guide* on Dell.com/Support/Manuals and on Dell.Avocent.com.

For additional detailed information on Dell's systems management portfolio, see the *Dell OpenManage Systems Management Overview Guide* on Dell.com/Support/Manuals.



Appendix A. Additional specifications

Enclosure dimensions and weight

Table 16 lists the dimensions and weight of the PowerEdge M1000e blade enclosure.

Table 16. Dimensions and weight

Dimension	Measurement
Width (not including rack ears)	447.5 mm
Height	440.5 mm
Depth (back of EIA flange to back of chassis)	753.6 mm
Total system depth (front bezel to PS latch)	835.99 mm

Environmental specifications

The PowerEdge M1000e and blade servers are compliant with Dell Environmental Specification 00109, which requires the servers to be able to withstand a rate of temperature change of 10°C/hr operational and 20°C/hr non-operational, including a maximum humidity rate of change for both operation and non-operation of 10% relative humidity per hour.

For more details, see the *Dell PowerEdge M1000e, M915, M910, M820, M710HD, M710, M620, M610x, M610, M520, and M420 Getting Started Guide* on Dell.com/Support/Manuals.

For additional information about environmental measurements for specific system configurations, see Dell.com/environmental_datasheets.

Video specifications

The iKVM supports a video display resolution range from 640x480 at 60Hz up to 1280 x 1024 x 65,000 colors (noninterlaced) at 75Hz.

Peripherals

Common peripherals for the M1000e include the following:

- An external USB DVD-ROM Drive, used for local installation of OS or other software
- A Dell 1U rack console for mounting a system administrator's control station directly into a Dell rack without sacrificing rack space needed for servers and other peripherals:
 - 17-inch LCD flat-panel monitor with height adjustment
 - Specially designed keyboard and trackball combination
 - Twin PS/2 connectors
 - SVGA video output
 - 1U rack-mounting kit
 - Simple installation
- Uninterruptible power supplies (UPS) for racks that provide temporary power source to bridge the critical moments after a power failure, allowing:
 - Time to save and back up the data being processed
 - Safely power down servers
 - Support for up to 5000 VA (3750 watts)



- Power Distribution Units — use the Dell Energy Smart Solution Advisor (ESSA) to see what a given chassis configuration requires.
 - Single-phase needs one power distribution unit per chassis
 - > 30A for a medium to lightly loaded chassis
 - > 60A for a heavily loaded chassis
 - For three-phase:
 - > 30A three-phase for a heavily loaded single chassis
 - > 50A or 60A three-phase for multiple chassis
 - Generally, two three-phase circuits are run to the rack through a power distribution unit, which then breaks out single phases to each power supply

Learn more at Dell.com/ESSA.



Appendix B. Standards compliance

The M1000e blade enclosure conforms to the industry standards in Table 17.

Table 17. Industry standard documents

Standard	URL for information and specifications
ACPI Advance Configuration and Power Interface Specification, v2.0c	acpi.info
Ethernet IEEE 802.3-2005	standards.ieee.org/getieee802/802.3.html
HDG Hardware Design Guide Version 3.0 for Microsoft Windows Server	microsoft.com/whdc/system/platform/pcdesign/desguide/serv_erdg.mspx
IPMI Intelligent Platform Management Interface, v2.0	intel.com/design/servers/ipmi
DDR3 Memory DDR3 SDRAM Specification, Rev. 3A	jedec.org/download/search/JESD79-3C.pdf
LPC Low Pin Count Interface Specification, Rev. 1.1	developer.intel.com/design/chipsets/industry/lpc.htm
PCI Express PCI Express Base Specification Rev. 2.0 and 3.0	pcsig.com/specifications/pciexpress
PMBus Power System Management Protocol Specification, v1.2	pmbus.info/specs.html
SAS Serial Attached SCSI, v1.1	t10.org
SATA Serial ATA Rev. 2.6; SATA II, SATA 1.0a Extensions, Rev. 1.2	sata-io.org
SMBIOS System Management BIOS Reference Specification, v2.7	dmtf.org/standards/smbios/
TPM Trusted Platform Module Specification, v1.2	trustedcomputinggroup.org
UEFI Unified Extensible Firmware Interface Specification, v2.1	uefi.org/specs
USB Universal Serial Bus Specification, Rev. 2.0	usb.org/developers/docs



Standard	URL for information and specifications
Windows Logo Windows Logo Program System and Device Requirements, v3.10	microsoft.com/whdc/winlogo/hwrequirements.mspx



Appendix C. Additional resources

Table 18 provides a list of documents and websites that provide for more information on the Dell PowerEdge M1000e.

Table 18. Additional resources

Resource	Description of contents	Location
Dell PowerEdge M1000e Enclosure Owner's Manual	This manual provides information on the PowerEdge M1000e chassis enclosure and its supported blade server modules.	Dell.com/Support/Manuals
Dell PowerEdge M1000e, M915, M910, M820, M710HD, M710, M620, M610x, M610, M520, and M420 Getting Started Guide	This guide is printed and shipped with the system, and is also available in PDF format on the Dell support site. This guide provides information on the following: <ul style="list-style-type: none"> • Initial setup steps • Key system features • Technical specifications 	Dell.com/Support/Manuals
Dell PowerEdge M1000e Systems Configuration Guide	This guide provides configuration information about the M1000e enclosure.	Dell.com/Support/Manuals
Dell Chassis Management Controller Firmware User's Guide	This guide provides information for installing, configuring, and using the Dell CMC firmware.	Dell.com/html/us/products/pedge/poweredge_mseries/blade_videos/poweredge.html
M1000e videos	Videos highlighting the major M1000e features.	On the module
System Information Label	The system information label documents the system board layout and system jumper settings.	On the module
PowerEdge Mxxx Technical Guide	These guides provide detailed technical information on the M820, M620, M520, or M420 supported features.	Dell.com/PowerEdge
Energy Smart Solution Advisor (ESSA)	The Dell online advisor console enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use ESSA to calculate the power consumption of your hardware, power infrastructure, and storage.	Dell.com/ESSA
Power and cooling technologies	Provides details for improving energy efficiency in the data center.	Dell.com/PNC
Energy management	Provides information on Dell's Fresh Air cooling solutions.	Dell.com/FreshAir



Resource	Description of contents	Location
Operating system matrix for Dell PowerEdge systems	Provides updated information on which operating systems are available on which PowerEdge systems.	Dell.com/OSsupport
Processor and chipset	Provides more information about server supported Intel processors and chipset.	Intel.com
Dell PowerEdge RAID controllers	Provides more information on Dell PowerEdge RAID controllers (PERC).	Dell.com/PERC
Power distribution unit (PDU)	Provides help selecting a PDU.	DellPDU.com
Uninterruptible power supply (UPS)	Provides help selecting a UPS model.	DellUPS.com
Volatility information	Contact your Dell Sales Representative.	Dell.com/Support/Manuals
Dell Enterprise Acoustics	White paper that explores the mechanisms of, people's reaction to, language of, and Dell's work to control noise from Enterprise products.	Dell.com/downloads/global/products/pedge/en/acoustical-education-dell-enterprise-white-paper.pdf
Dell Regulatory Certification	Product Safety, EMC, and Environmental Datasheets	Dell.com/regulatory_compliance_datasheets

