


PowerEdge M1000e



Technical Guide



The M1000e chassis provides flexibility, power and thermal efficiency with scalability for future needs.

Dell

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, PowerEdge, PowerConnect, RapidRails, VersaRails, FlexAddress, and OpenManage are trademarks of Dell, Inc. Avocent is a registered trademark of Avocent Corporation or its subsidiaries. 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. Brocade is a registered trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. Cisco and Catalyst are registered trademarks of Cisco and/or its affiliates in the U.S. and certain other countries. Citrix® and XenServer™ are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. Egenera registered trademarks of Egenera, Inc. in the United States and/or other countries. InfiniBand is a registered trademark and service mark of the InfiniBand Trade Association. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Mellanox is a registered trademark of Mellanox Technologies, Inc. Microsoft, Windows Server, Active Directory, and Hyper-V are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Platespin and PowerConvert are registered trademarks of Novell, Inc., in the United States and other countries. Velcro is a registered trademark of Velcro Industries B.V. VMware is a registered trademark and vCenter is a trademark of VMware, Inc. in the United States and/or other jurisdictions. Dell disclaims proprietary interest in the marks and names of others.

©Copyright 2010 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.

Initial Release June 2010

Table of Contents

1	Product Comparison	5
2	New Technologies.....	7
2.1	Overview	7
2.2	Detailed Information	7
3	System Information	9
3.1	Overview	9
3.2	Product Features Summary	9
4	Mechanical	10
4.1	Chassis Description.....	10
4.2	Dimensions and Weight.....	10
4.3	Front Panel View and Features	10
4.4	Back Panel Features	12
4.5	Power Supply Indicators.....	12
4.6	Rails and Cable Management	13
4.7	Rack Support.....	16
4.8	Rack View	16
4.9	Fans	17
4.10	Cabling.....	21
4.11	Control Panel/LCD	22
4.12	Security.....	24
5	Power, Thermal, Acoustic	26
5.1	Power Supplies	26
5.1.1	Supported Voltages	27
5.1.2	Redundancy	27
5.1.3	Power Management	29
5.2	Power Supply Specifications.....	30
5.3	Heat Dissipation	30
5.4	Environmental Specifications.....	33
5.5	Power Consumption.....	33
5.6	Maximum Input Amps.....	33
5.7	Power-Up Sequence	33
5.8	Acoustics	33
6	Processors and Memory	35
7	Midplane.....	36
8	Embedded NICs/LAN on Motherboard (LOM)	39
9	I/O	40
9.1	Overview	40
9.2	Quantities and Priorities	41
9.3	Supported Mezzanine Cards and Switches.....	44
9.4	I/O Module Installation.....	45
9.5	FlexAddress	45
10	Storage	51
11	Video.....	52
12	Rack Information.....	53

Dell

12.1	Overview	53
12.2	Rails	53
12.3	Cable Management Arm (CMA)	54
12.4	Rack View	55
13	Virtualization	57
14	Systems Management.....	59
14.1	Overview	59
14.2	Server Management.....	60
14.3	Enclosure Management.....	61
14.4	Integrated Keyboard and Mouse Controller (iKVM).....	65
15	Peripherals.....	68
16	Packaging Options	69

Tables

Table 1.	Comparison of PowerEdge 1855/1955 Chassis and M1000e Chassis	5
Table 2.	Rack vs. Blade Server Rack-Level Specification Comparison	6
Table 3.	Feature Summary	9
Table 4.	Dimensions	10
Table 5.	Typical Modular Server System Rack Height and Cable Reduction.....	13
Table 6.	Fabric Specifications.....	43
Table 7.	FlexAddress Features and Benefits	47

Figures

Figure 1.	Server Density Comparison	5
Figure 2.	M1000e Front View.....	10
Figure 3.	Possible Server Module Sizes, Front Panel View	11
Figure 4.	Example Server Module Configurations	11
Figure 5.	Power Supply Indicators.....	12
Figure 6.	Rack Cabling	14
Figure 7.	RapidRails Rack Kit Contents	15
Figure 8.	VersaRails Rack Kit Contents	15
Figure 9.	M1000e in a Rack.....	16
Figure 10.	Rear View Showing Fans.....	17
Figure 11.	Blades, Blanks, and 1 Open Slot Needing to be Filled	18
Figure 12.	Power Supply, Power Supply Blanks, and Open Slot Needing to be Filled	18
Figure 13.	I/O Module and Open Slot Needing to be Filled	19
Figure 14.	Installed CMC, I/O Module, and Power Supply Blanks	20
Figure 15.	Installed iKVM Blank	20
Figure 16.	Power Supply, CMC, and I/O Module Blanks.....	21
Figure 17.	Simplified Cabling.....	22
Figure 18.	M1000e LCD Panel Recessed Position	23
Figure 19.	M1000e LCD Panel During Usage	23
Figure 20.	LCD Panel Capabilities.....	24
Figure 21.	Power Supplies in M1000e.....	26
Figure 22.	M1000e Power Supply Rear View.....	27

Dell

Figure 23.	Power Architecture	28
Figure 24.	PMBus Communication Channels	30
Figure 25.	Server Cooling Air Profile	31
Figure 26.	I/O Module Inlet and IOM Locations	31
Figure 27.	I/O Cooling Air Profile	32
Figure 28.	Power Supply Inlet and Cooling Air Profile	32
Figure 29.	Midplane	36
Figure 30.	M1000e Midplane Front View	37
Figure 31.	M1000e Midplane Rear View	38
Figure 32.	M1000e I/O Modules	40
Figure 33.	High Speed I/O Architecture	42
Figure 34.	Ethernet Growth Path	43
Figure 35.	Difference Between Passthroughs and Switch Modules	44
Figure 36.	FlexAddress Addresses	46
Figure 37.	FlexAddress Screen in the CMC	47
Figure 38.	FlexAddress SD Card	48
Figure 39.	SD Slot on bottom of CMC	48
Figure 40.	CMC FlexAddress Summary Screen	49
Figure 41.	CMC FlexAddress Server Detail Screen	50
Figure 42.	Examples of Major Storage Platforms Supported	51
Figure 43.	M1000e RapidRails Static Rails	53
Figure 44.	M1000e VersaRails Static Rails	54
Figure 45.	M1000e Strain Relief Bar and Cable Enumerator Clip (12 Per Kit)	55
Figure 46.	M1000e Mounted in the Rack	55
Figure 47.	M1000e Strain Relief Bar and Cable Enumerator Clips	56
Figure 48.	Examples of Major Virtualization Platforms Supported	57
Figure 49.	Examples of I/O modules Recommended for Use in Virtualized Environments	58
Figure 50.	System Management Architecture Simplified Block Diagram	60
Figure 51.	Chassis Management Controller	63
Figure 52.	CMC Module Features	64
Figure 53.	M1000e iKVM	65
Figure 54.	Rear iKVM interface Panel	66
Figure 55.	Front Keyboard/Video Ports	66
Figure 56.	Enclosure After Unpacking	69

1 Product Comparison

The Dell™ PowerEdge™ M1000e offers significant enhancements over its predecessor, the 1955, as can be seen in the following table:

Table 1. Comparison of PowerEdge 1855/1955 Chassis and M1000e Chassis

Feature	1855/1955 Chassis	M1000e Chassis
Blade Compatibility	PowerEdge 1855/1955	PowerEdge M600/M605 11G and beyond
Form Factor	7U	10U
No. of Blades	10	16
I/O Module Bay	4	6
Fabric Types Supported	1 x Dual GbE 1 x Dual Xaui 1 Lane - GbE, FC2 4 Lane - 4 x IB	2 x 2 Lane to support: GbE2 x 4 2 X 4 Lane to support: 1 Lane - GbE, 10GbE serial/KR, FC8/4/2/1 4 Lane - IB, 10GbE (Xaui. KR), 40GbE
Power Supplies	2 x (non-redundant) or 4 x 2100W PSUs	3 x non-redundant) or 6 x 2360W PSUs
Management Modules	1 (std) 2nd(optional)	1 (std) 2nd(optional)
KVM options	1 x Avocent® Analog or Digital KVM	1 x Avocent® Analog KVM (optional)

Putting 16 half-height blades in the PowerEdge M1000e is 60% more dense than using 1U servers.




Figure 1. Server Density Comparison

Dell

Greater density means:

- Smaller Footprint
- More Processing Performance
- More RAM capacity
- Lower Power Consumption per unit
- Easier Manageability

Dell’s blade server platform offers superior feature density over comparable rack servers, as can be seen from 0. (Darker blue shading indicates increased memory density.)

Table 2. Rack vs. Blade Server Rack-Level Specification Comparison¹

	R410	R510	R610	R710	R810	R815	R905	R910	M605	M610	M710	M805	M905	M910
Form Factor	Rack ²	Rack	Rack	Rack	Rack	Rack	Rack	Rack	1/2 Blade	1/2 Blade	Full Blade	Full Blade	Full Blade	Full Blade
Processors														
Manufacturer	Intel	Intel	Intel	Intel	Intel	AMD	AMD	Intel	AMD	Intel	Intel	AMD	AMD	Intel
Sockets	2	2	2	2	4	4	4	4	2	2	2	2	4	4
Max Cores per 42U Rack	504	252	504	252	672	1,008	240	320	768	768	384	384	768	1,024
Memory														
Max RAM per rack, in TB	5	3	8	4	11	5	3	10	4	12	6	4	6	16
I/O														
Max 1GB Ethernet Ports per 42U Rack	252	378	504	420	588	588	320	440	512	640	576	512	512	640
Max 10GbE, DDR IB, or FC network ports per rack	84	168	168	84	252	252	140	220	256	256	256	256	256	256
Internal Storage³														
Drives per 42U rack	168	252	252	168	126	126	50	160	128	128	128	64	64	64
Max 7.2k or 10k rpm internal storage per rack	336	504	151	252	76	76	50	96	77	77	77	38	38	38
Max 15k rpm internal storage per 42U rack	101	151	37	76	18	18	15	23	19	19	19	9	9	9
Max SSD internal storage per 42U rack	17	25	25	17	13	13	0	16	13	13	13	6	6	6

¹ This rack-level physical capacity specification summary does not factor in power and cooling.

² 42U is the most common rack size.

³ Storage measurements provided in Terabytes.

2 New Technologies

2.1 Overview

The PowerEdge M1000e is designed to help customers 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, 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 total cost of ownership (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 passthrough modules), a high-performance and highly available passive midplane that connects server modules to the infrastructure components, power supplies, fans, integrated KVM and Chassis Management Controllers (CMC). The PowerEdge M1000e uses redundant and hot-pluggable components throughout to provide maximum uptime.

The M1000e provides identical and symmetric 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, 10GE copper ports, or 10GE optical ports. True modularity at the system and subsystem level provides simplicity of extension and enhancement, now and in the future.

The main benefits to customers of these features include improved:

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

Together, these factors enable customers to do more with their server investment.

2.2 Detailed Information

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

The chassis features:

Dell

- A high-speed passive midplane that connects the server modules in the front and power, I/O, and management infrastructure in the rear of the enclosure.
- Comprehensive I/O options to support dual links of 40 Gigabits per second today (with 4x QDR InfiniBand®) with future support of even higher bandwidth I/O devices when those technologies become available. This provides 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 Chassis Management Controller (CMC), Keyboard/Video/Mouse (KVM) switch, and server modules.
- Up to two Chassis Management Controllers (CMC-1 is standard, CMC-2 provides optional redundancy) and 1 optional integrated Keyboard/Video/Mouse (iKVM) switch.
- Up to 6 hot-pluggable, redundant Power Supplies and 9 hot-pluggable, N+1 redundant fan modules.
- System Front Control panel w/ LCD panel and two USB Keyboard/Mouse and one Video “crash cart” connections.

3 System Information

3.1 Overview

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.

3.2 Product Features Summary

Table 3. Feature Summary

Feature	Parameter
Chassis Size	10U high rack mount
Blades per Chassis	16 Half Height, 8 Full Height
Total Blades in a 42U Rack	64 Half Height, 32 Full Height
Total I/O Module Bays	6 (3 redundant or dual fabrics)
Total Power Supplies	6 (3+3 redundant)
Total Fan Modules	9 (8+1 redundant)
Management Modules and Interfaces	2 CMCs (1+1 redundant), 1 iKVM, Front Control Panel, Graphical LCD Control Panel
AC Redundancy	3+3 2+2 1+1 Each requires power supplies in slots 1, 2, and 3 to be connected to a different grid as compared to those in slots 4, 5, and 6.
DC Redundancy	1+1 2+1 3+1 4+1 5+1 Each with one extra power supply that comes online if one of the existing power supplies fails.

4 Mechanical

4.1 Chassis Description

The Dell M1000e supports up to sixteen 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.

4.2 Dimensions and Weight

Table 4. Dimensions

Dimension	Measurement
Width, not including rack ears	447.5 mm
Height	440.5 mm
Depth, Rear of EIA Flange to Rear of Chassis	753.6 mm
Total System Depth (Front Bezel to PS Latch)	835.99 mm

4.3 Front Panel View and Features




Figure 2. M1000e Front View

The M1000e enclosure supports up to 16 half-height or 8 full-height server modules, each occupying a slot accessible in the front of the enclosure. The enclosure has also been designed to accommodate other form factors, including dual-width modules.

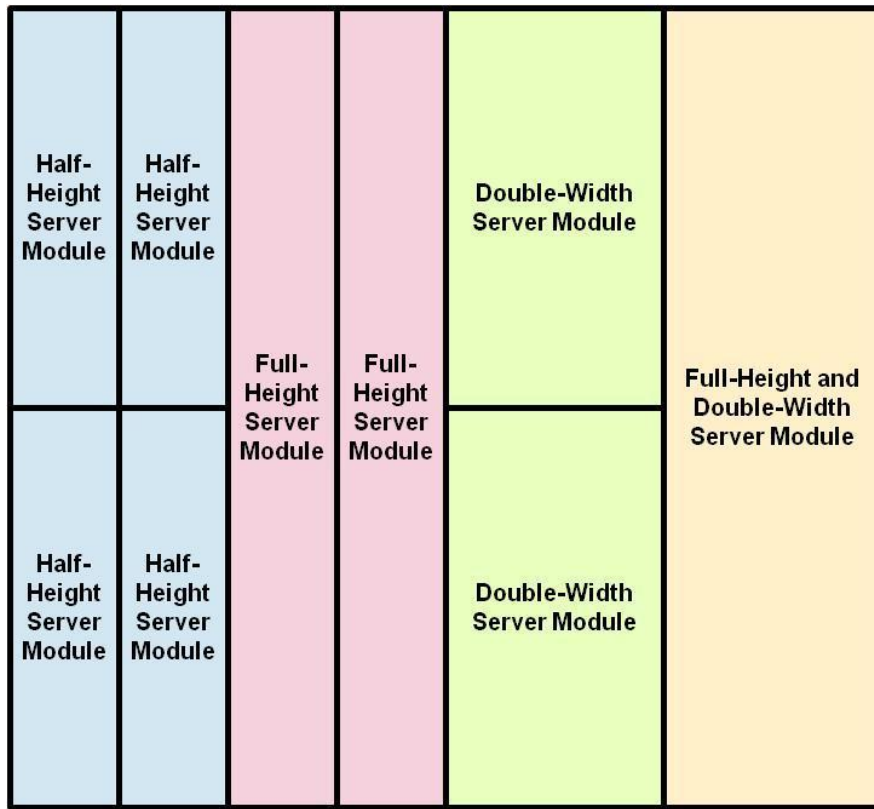


Figure 3. Possible Server Module Sizes, Front Panel View

Server Modules can be freely located within each 2 x 2 half-height quadrant. The mechanical design of the M1000e has support structures for half-height server modules above or below double-width server modules, and for half-height server modules side-by-side with full-height server modules.

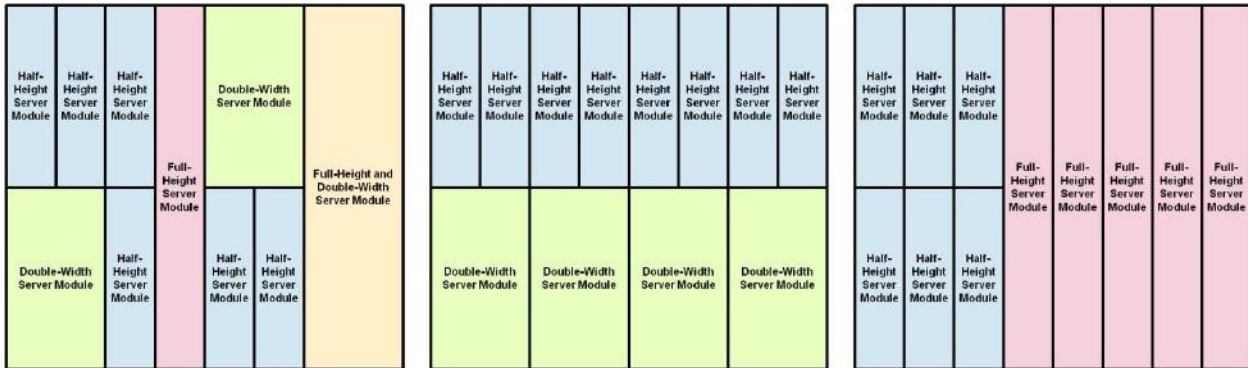


Figure 4. Example Server Module Configurations

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 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. For more information, see [System Control Panel Features](#) in the *Hardware Owner's Manual*.

Dell

Fresh air plenums are at both 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.

4.4 Back Panel Features

The rear of the M1000e Enclosure contains system management, cooling, power and I/O components. At the top of the enclosure are slots for two Chassis Management 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. For more information, see [Back-Panel Features](#) in the *Hardware Owner's Manual*.

4.5 Power Supply Indicators

Figure 5 shows the power supply indicators. For more information, see [Back-Panel Features](#) in the *Hardware Owner's Manual*.




Figure 5. Power Supply Indicators

4.6 Rails and Cable Management

RapidRails™ Static Rails for Square Hole Racks:

- Supports toolless installation in 19” EIA-310-E compliant square hole 4-post racks including all generations of Dell racks except for the 4200 & 2400 series
- Has a minimum rail depth of 703 mm
- Provides a square-hole rack adjustment range of 712-755 mm
- Includes strain relief bar and cable enumerators for managing and securing cables

VersaRails™ Static Rails for Square or Round Hole Racks:

- Supports tooled installation in 19” EIA-310-E compliant square or unthreaded round hole 4-post racks
- Has a minimum rail depth of 703 mm
- Provides a square-hole rack adjustment range of 706-755 mm
- Provides a round-hole rack adjustment range of 706-755 mm
- Includes strain relief bar and cable enumerators for managing and securing cables

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. The following table 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 the table 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 system installation in Dell or other industry-standard racks.

Table 5. 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	4x4	-	-
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




Figure 6. Rack Cabling

RapidRails™ Static Rails for Square Hole Racks supports toolless installation in 19" EIA-310-E compliant square hole 4-post racks including all generations of Dell racks except for the 4200 & 2400 series. Minimum rail depth is 703 mm. Square-hole rack adjustment range is 712-755 mm. The rail system includes a strain relief bar and cable enumerators for managing and securing cables.

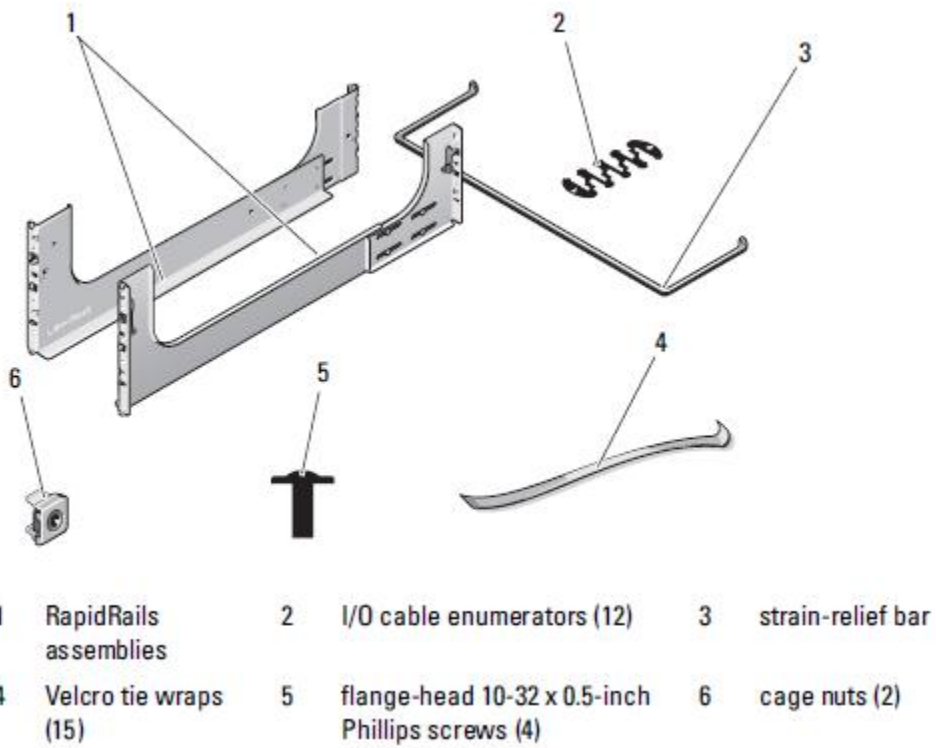


Figure 7. RapidRails Rack Kit Contents




Figure 8. VersaRails Rack Kit Contents

See Section 12 for more details.

Dell

4.7 Rack Support

The M1000e chassis offers the following options for rack support:

- RapidRails™ static rails for toolless mounting in 4-post racks with square holes
- VersaRails™ static rails for tooled mounting in 4-post racks with square or unthreaded round holes

See Section 12 for more details.

4.8 Rack View



Figure 9. M1000e in a Rack

4.9 Fans



Figure 10. Rear View Showing Fans

The PowerEdge M1000e chassis comes standard with 9 hot-swappable, redundant fan modules that are distributed evenly across the enclosure. The speed of each fan is individually managed by the CMC. Together, these design innovations can provide:

- Significant power savings as compared to older servers
- Less airflow required as compared to the same number of similarly configured 1U 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 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 on-board monitoring. Failure of more than one fan will not automatically result in shutting down of blade servers. This is because the blade servers 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/airflow reasons.



Figure 11. Blades, Blanks, and 1 Open Slot Needing to be Filled




Figure 12. Power Supply, Power Supply Blanks, and Open Slot Needing to be Filled



Figure 13. I/O Module and Open Slot Needing to be Filled



Figure 14. Installed CMC, I/O Module, and Power Supply Blanks



Figure 15. Installed iKVM Blank



Figure 16. Power Supply, CMC, and I/O Module Blanks

4.10 Cabling

There are two types of external cabling simplification features offered:

- Stacked Ethernet Switching
 - Internal switches have optional 10GbE uplinks and/or stacking connectors
 - Manage/configure multiple switches as one with stacking
 - Consolidate uplinks from multiple chassis into 2-4 x 10GbE ports
- Stacked CMCs
 - CMC has a 2nd 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 17. Simplified Cabling

4.11 Control Panel/LCD

The control panel contains the local user interface. Functions include chassis level diagnostic LEDs, LCD panel, and power button. This device is hot-pluggable and is always powered, even in chassis standby mode.



Figure 18. M1000e LCD Panel Recessed Position



Figure 19. M1000e LCD Panel During Usage

The M1000e chassis LCD shows extensive information about the status of each hardware module, network information for the CMC and each iDRAC, and status messages with detailed explanations in plain language. Users may access a wide variety of information about modules via the panel, including their type, user-defined name, configurations, service tag numbers, and IP address

Dell

information. The LCD panel can be retracted into the chassis body, or extended and angled once deployed for full visibility no matter where the M1000e is mounted in the rack.

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 network configuration.

Figure 20 shows some of the capabilities of the LCD control panel.

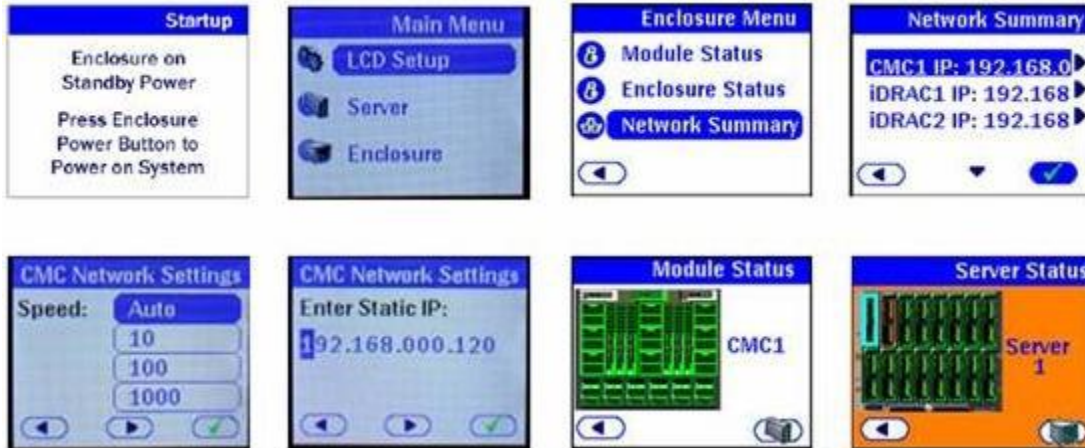


Figure 20. LCD Panel Capabilities

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 in each blade
- Status information screens for each blade
- Status information screens for the modules installed in the back of the enclosure, including the IO 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

See the [M1000e Configuration Guide](#) and the [CMC Administrator Reference Guide](#) for more details on the capabilities of the LCD panel.

4.12 Security

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

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

Dell

Most of the security capabilities are driven by the CMC, which provides a mechanism for centralized configuration of the M1000e enclosure's security settings and user access. It 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 the Web interface
- Web interface supports 128-bit SSL 3.0 encryption and 40-bit SSL 3.0 encryption (for countries where 128-bit is not acceptable)
- 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
- Disabling front panel access

5 Power, Thermal, Acoustic

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, increasing efficiencies of DC-DC converters, and replacing some linear voltage regulators with more-efficient switching regulators.

See Section 15 for more information on external power connection accessories.

5.1 Power Supplies

The power distribution inside the M1000e Modular Server System consists of a power supply system located in the rear bottom of the chassis.



Figure 21. Power Supplies in M1000e



Figure 22. M1000e Power Supply Rear View

The Dell power supplies utilize output Oring FETs to isolate the power supply from the 12V system bus. If a single power supply fails its output Oring FET for that power supply will turn off removing itself from the bus. Think of it as an electrical switch that turns off when the power supply fails.

5.1.1 Supported Voltages

Dell currently offers a power supply rated at 2360W 230V. With current sharing between power supplies, total system redundant power is approximately 7080W in a 3+3 power supply configuration.

5.1.2 Redundancy

Power redundancy in the M1000e supports any necessary usage model, though the M1000e requires three power supplies to power a fully populated system or six power supplies in a fully redundant system.

AC redundancy is supported in the following configurations, each of which requires the power supplies in slots 1, 2, and 3 to be connected to a different grid as compared to those in slots 4, 5, and 6.

- 3+3
- 2+2
- 1+1

DC redundancy is supported in the following configurations, each with one extra power supply that comes online if one of the existing power supplies fails:

Dell

- 1+1
- 2+1
- 3+1
- 4+1
- 5+1

When Dynamic Power Supply Engagement (DPSE) is enabled, the PSU units move between On and Off states depending upon actual power draw conditions to achieve high power efficiency by driving fewer supplies to maximum versus all with partial and less-efficient loading.

In the N+N power supply configuration, the system will provide 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 12 Volt DC power for the system, both main power and standby power. The CMCs, LCD and Control Panel are powered solely by 12 Volt Standby power, insuring 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 12 Volt Main power.

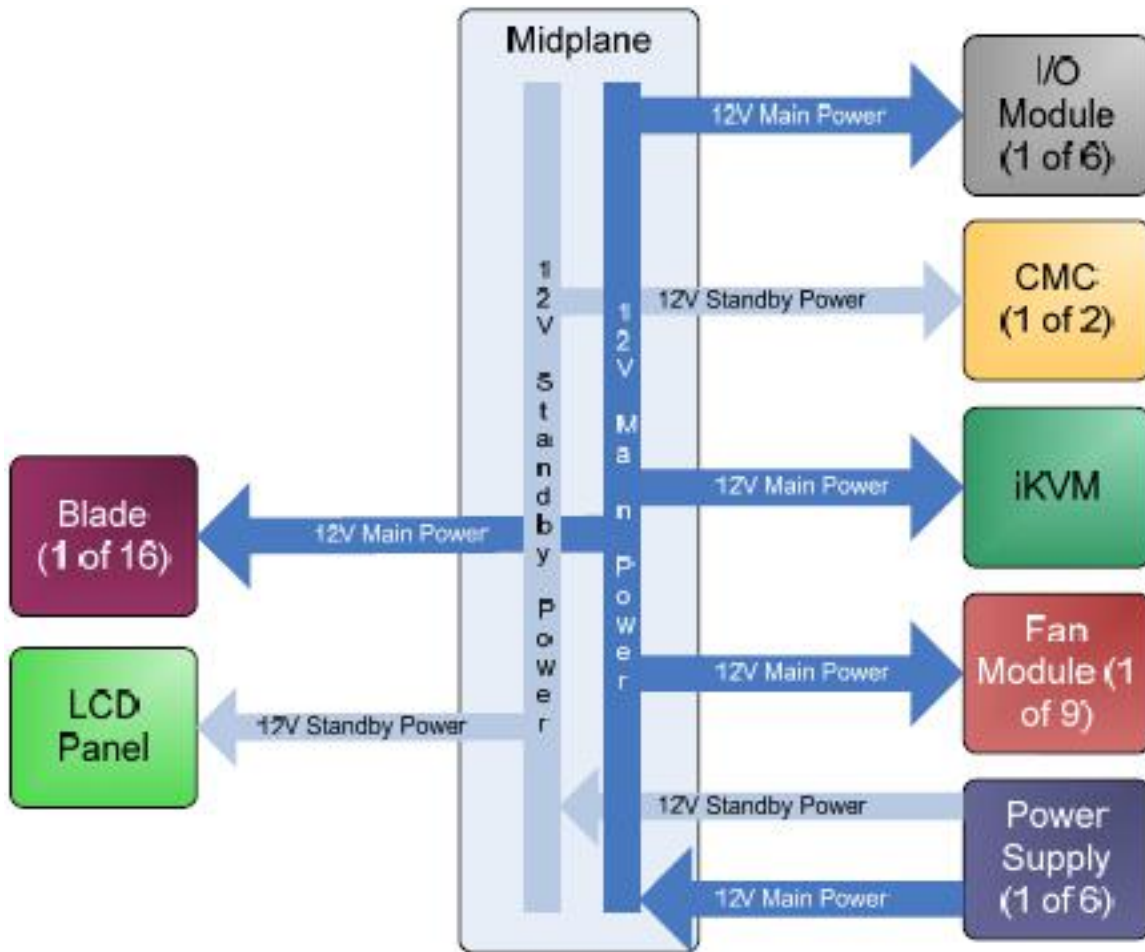


Figure 23. Power Architecture

5.1.3 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 to the user, 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 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 performs a sophisticated power budget inventory for the server module, based upon its configuration of CPUs, memory, I/O and local storage. Once this number is generated, the iDRAC 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 hardware constantly monitors actual power consumption at each server module. This power measurement is used locally by the server module to insure 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 regards 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 insure 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 processor, memory can throttle down when necessary on lower priority blade servers. An allocation is taken out for the infrastructure (fans, IO modules) and then the remainder is applied to the blades, and then throttling is applied if required to get under the cap. If all the blades are setup with the same priority, then it will start throttling down processor, memory, and so on. A variety of BIOS settings will 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 will not shut down; rather they will 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, will 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 Dell™ OpenManage™ IT Assistant or to other enterprise-level management tools. Real time power consumption is viewable per system.

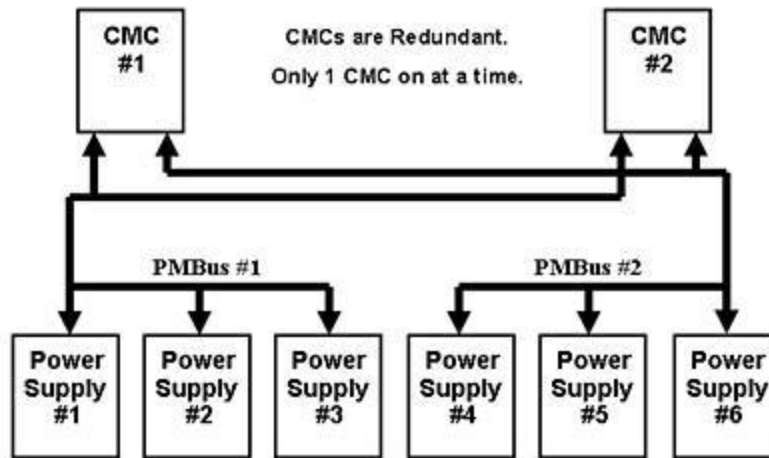


Figure 24. PMBus Communication Channels

All VMware® products include consuming the "current power consumption" and "current power cap/limit" retrieval via Dell specific IPMI commands using 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.

The Power Management chapter in the [Dell Chassis Management Controller \(CMC\) User Guide](#) provides extensive information on power management.

5.2 Power Supply Specifications

Each power supply offers:

- 91%+ AC/DC Conversion Efficiency
- Dynamic Power Supply Engagement which automatically engages the minimum number of supplies required to power a given configuration, maximizing power supply efficiency

The following are the PowerEdge M-1000e chassis power supply capabilities:

- 2360 watts maximum for each PSU
- 220 VAC (Volts Alternate Current) input (a single PSU runs between 180V and 260V AC)
- 50Hz or 60Hz input
- 14A maximum input
- 192A (Amps) @ + 12 Volts DC (Direct Current) output Operational
- 4.5A @ +12 Volt output Standby
- 3 or 6 PSU configurations available
- PSUs are hot-swappable

5.3 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

Dell

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 in the customer's required power-to-cool.

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.

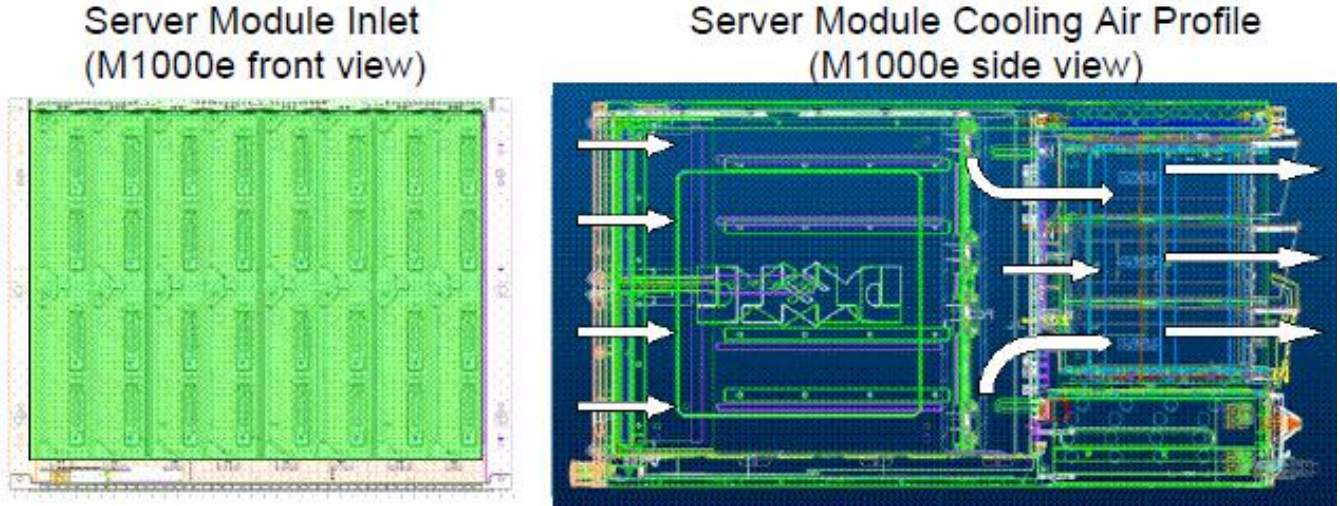


Figure 25. Server Cooling Air Profile

The Server Modules are cooled with traditional front-to-back cooling. As shown in the figure, 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 which 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.

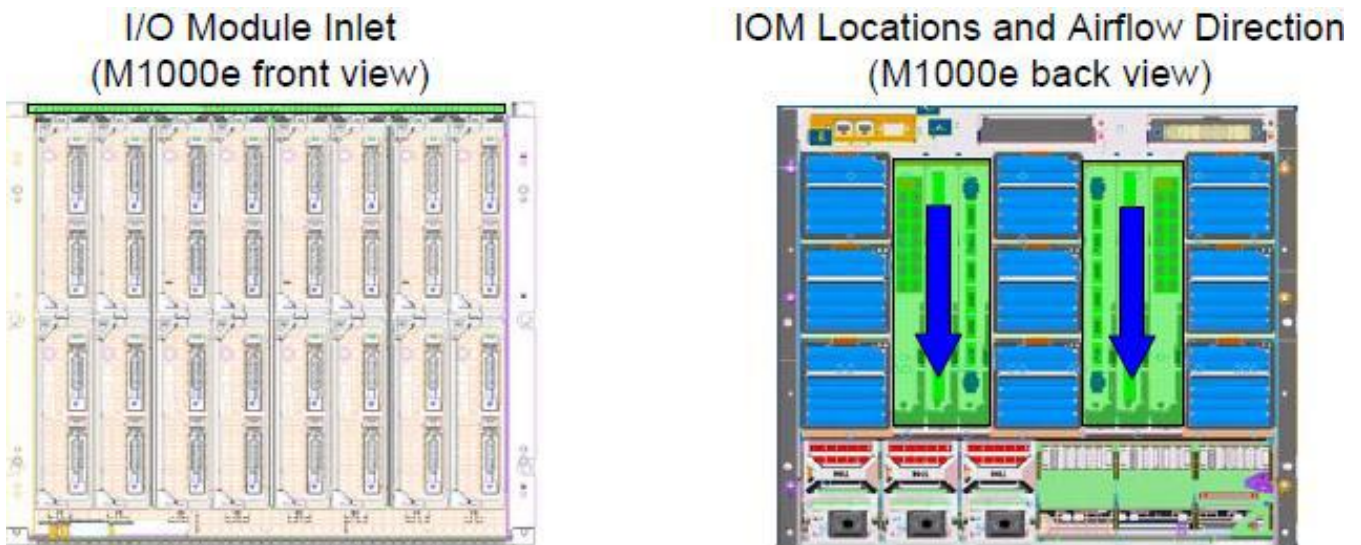


Figure 26. I/O Module Inlet and IOM Locations

I/O Cooling Air Profile (M1000e side view)




Figure 27. I/O 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 the figure. 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.

Power Supply Inlet (M1000e front view)



Power Supply Cooling Air Profile (M1000e side view)




Figure 28. Power Supply Inlet and Cooling Air Profile

The Power Supplies, located in the rear 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 the figure above. This insures that the power supplies receive ambient temperature air.

Dell

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

- Server module level thermal monitoring by the iDRAC
- 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.

5.4 Environmental Specifications

See the [Getting Started Guide](#) on support.dell.com.

5.5 Power Consumption

Use the [Dell Energy Smart Solution Advisor](#) (ESSA) to see requirements for a specific chassis configuration.

5.6 Maximum Input Amps

See [Power Distribution Systems for the Dell M1000e Modular Server Enclosure - Selection and Installation](#).

5.7 Power-Up Sequence

The following steps detail how and in what order the M1000e components are powered up:

1. The first power supply provides a small amount of electricity which starts up the first CMC.
2. CMC begins to boot and power up the power supply units.
3. Active and Standby CMC boot up Linux® operating system.
4. Active CMC powers up all remaining PSUs.
5. All six PSUs are powered up.
6. Server iDRAC are powered up. (In slot priority order from 1-9; i.e., highest priority 1 slots first, then priority 2, etc. If all same priority, goes in slot order 1-16. Each one is spaced apart by 500ms.)
7. iKVM is powered up.
8. IOM modules are powered up.
9. Depending upon Blade BIOS power setting (last power state, always on or always off), blade iDRAC requests power up from CMC first come/first served in order from #6.
10. CMC powers up blades.

For full configuration, booting the enclosure takes between 2-4 minutes, followed by 1-4 minutes for each blade.

5.8 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

Dell

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 ESSA tool: <http://solutions.dell.com/DellStarOnline/Launch.aspx/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. Please ensure that components are operating properly if fans remain at full speed.
- The CMC will automatically raise and lower the fan speed to a setting that is appropriate to keep all modules cool.
- If a single fan is removed, all fans will be 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 will cause the rest of the fans to settle back to a quieter state.
- Whenever communication to the CMC or iDRAC is lost such as during firmware update, the fan speed will increase and create more noise.

6 Processors and Memory

With the addition of the PowerEdge M910 server to the PowerEdge portfolio, the M1000e is now scalable to 256 cores & 4TB of RAM: 4 sockets x 8 cores x 8 blades = 256 Cores; 32 DIMM sockets x 16GB DIMMs x 8 Blades = 4096GB or 4TB RAM.

See the *Technical Guide* for each of the compatible blade servers offered for more details on processors and memory offered.

7 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 29. 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.

The midplane serves as transport for a patent-pending, time-division-multiplexed serial bus for general purpose I/O reduction. This serial bus contributes greatly to the midplane's I/O lane count reduction, which is typically burdened with a significant I/O pin and routing channel count of largely static or low-speed functions. For instance, all Fibre Channel I/O Passthrough module LED and SFP status information is carried over this bus, which alone eliminates over one hundred point-to-point connections that would otherwise be required. The time division multiplexed serial bus is fully redundant, with health monitoring, separate links per CMC and error checking across all data.

Blade
Signal
Connectors

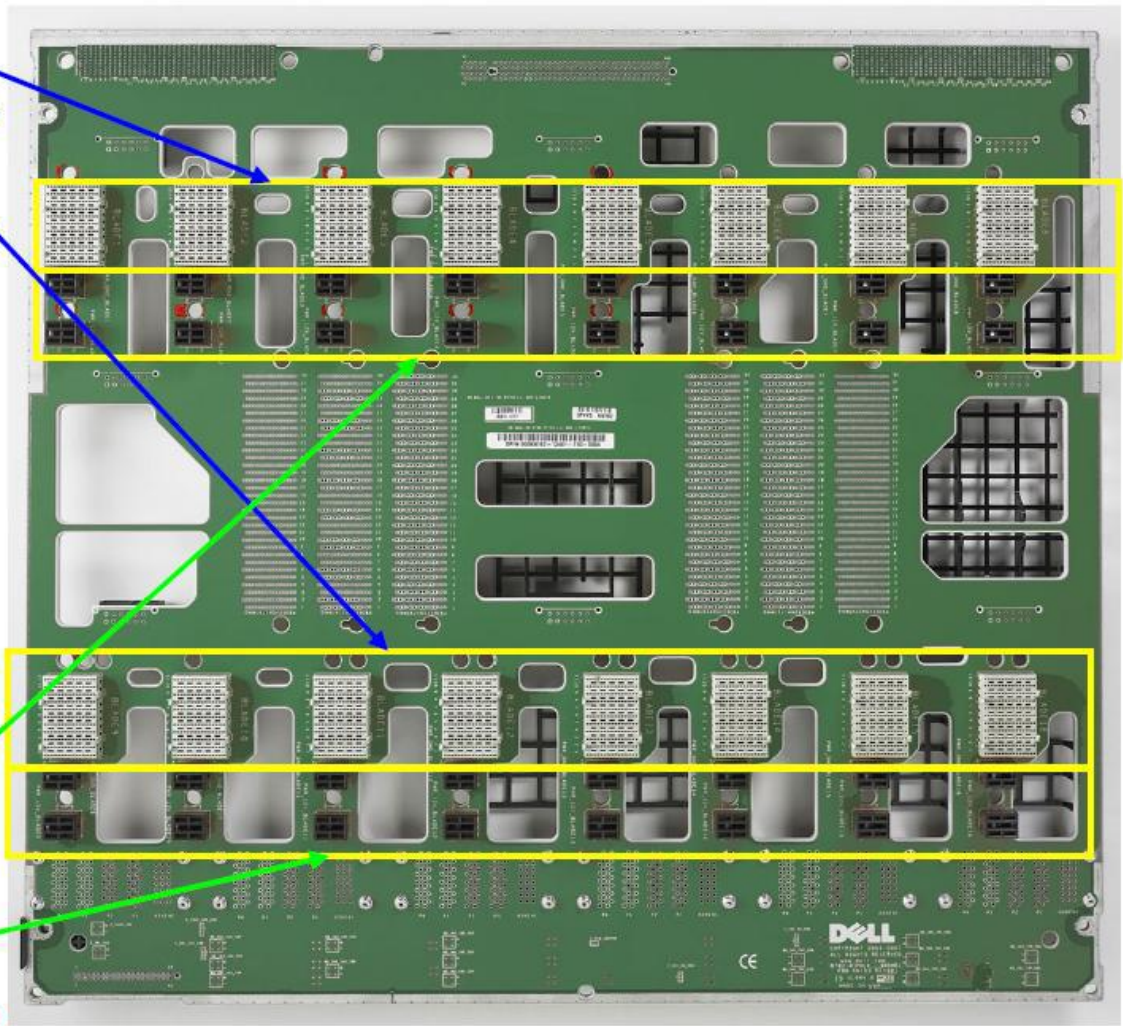


Figure 30. M1000e Midplane Front View

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 3 axes. This provides close tolerance alignment between the server modules and their midplane connections. The midplane has been carefully designed to minimize the impact to the overall system airflow.

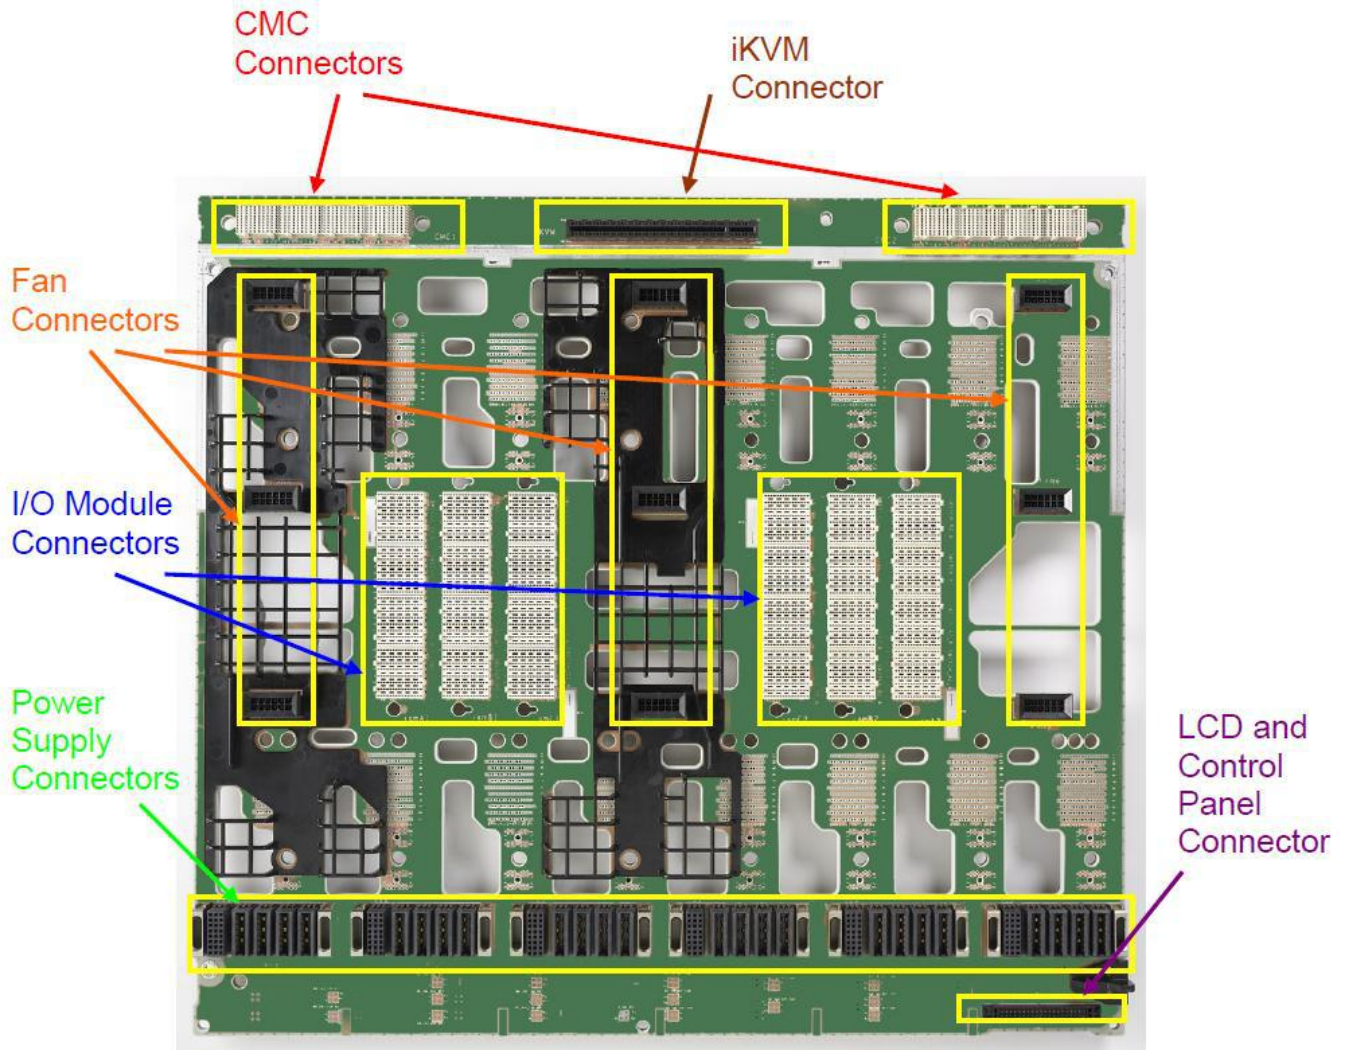


Figure 31. M1000e Midplane Rear View

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

Dell

8 Embedded NICs/LAN on Motherboard (LOM)

See the *Technical Guide* for each of the compatible blade servers.

9.1 Overview

Dell M-series provides complete, snap-in FlexI/O scalability down to the switch interconnects. Flex I/O technology is the foundation of the M1000e I/O subsystem. Customers may mix and match I/O modules, including Cisco®, Dell™ PowerConnect™, Fibre Channel, and InfiniBand options. The I/O modules may be installed singly or in redundant pairs. See [I/O Connectivity](#) in the *Hardware Owner's Manual* for detailed information.




Figure 32. M1000e I/O Modules

These I/O modules are connected 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 customers 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
- Quad Data Rate InfiniBand Switch options available for HPCC
- Up to 8 high-speed ports

Dell

- 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/EMC SAN fabrics
- High-availability clustering inside a single enclosure or between two enclosures

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).

9.2 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 via the midplane to the associated I/O Modules (IOM) in the rear of the enclosure, which then connect to the customer's LAN/SAN/IPC networks.

The optional mezzanine cards are designed to connect via 8-lane PCIe to the server module's chipset in most cases. Mezzanine cards may have either one dual port ASIC with 4- or 8-lane PCIe interfaces or dual ASICs, each with 4-lane PCIe interfaces. External fabrics are routed through high-speed, 10-Gigabit-per-second-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 the accidental activation of any misconfigured fabric device on a server module. The system will automatically detect this misconfiguration and alert the user of the error. No damage occurs to the system, and the user will have the ability to 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 33.

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 passthrough modules enabling you to connect to any legacy infrastructure whether using Ethernet passthrough or switch technology. This technical advance uses in-band signaling on 1000BASE-KX transport and requires no user interaction for enablement.

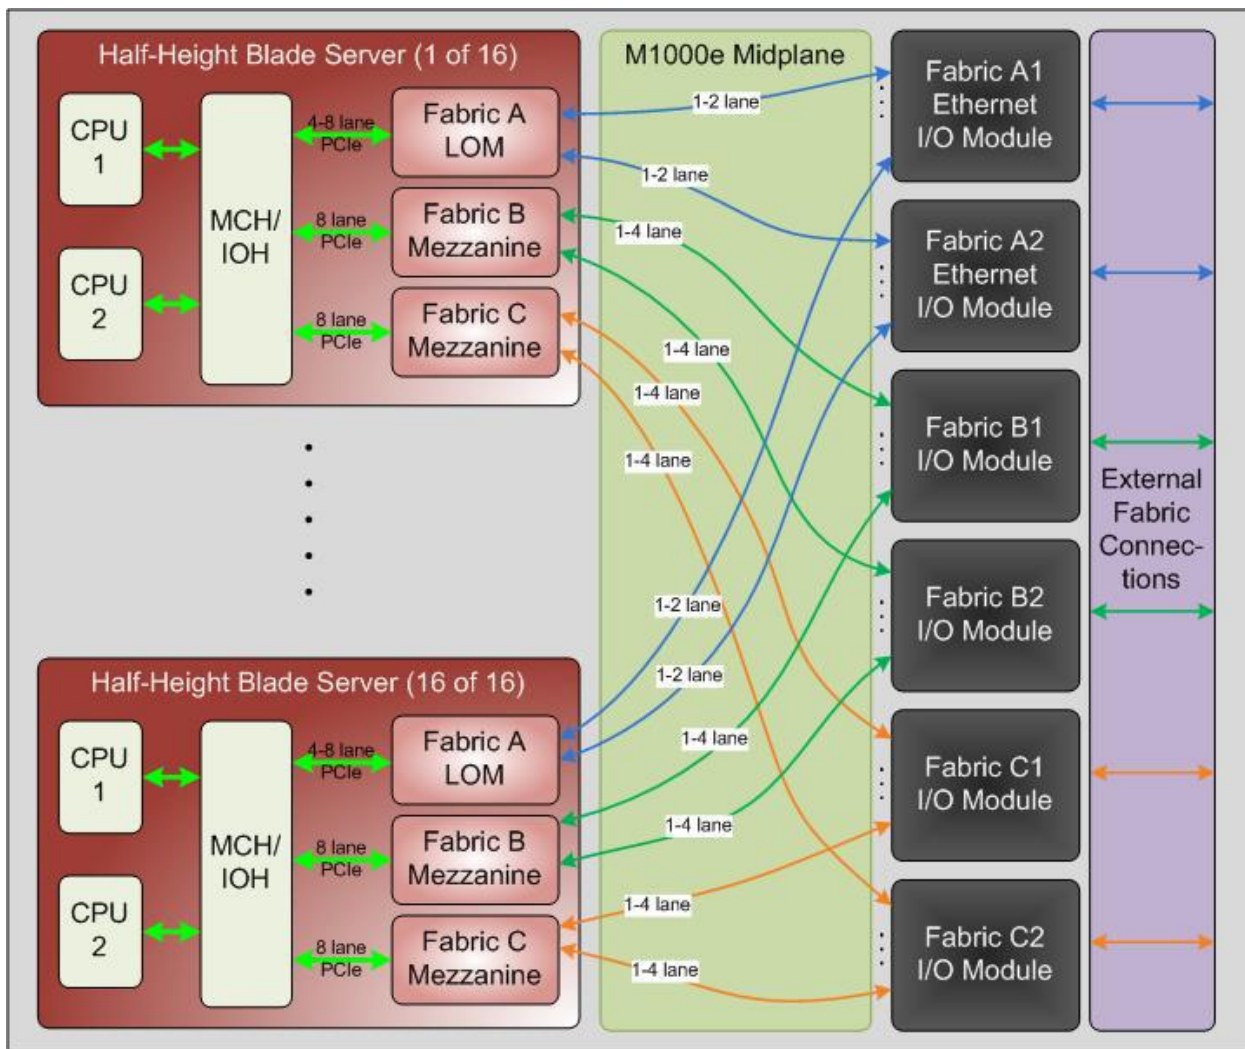


Figure 33. High Speed I/O Architecture

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 rear of the chassis. Supported bandwidth ranges from 1 to 10 Gbps per lane depending on the fabric type used.

Table 6. 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 Gen1	8B/10B	2.5	2	8 (4 lane)	1,2,4,8,12,16,32
PCIe Gen2	8B/10B	5	4	16 (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 8bps	8B/10B	8.5	6.8	6.8	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
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




Figure 34. Ethernet Growth Path

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 passthrough 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 passthroughs or switches. Passthrough 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 customer's network.

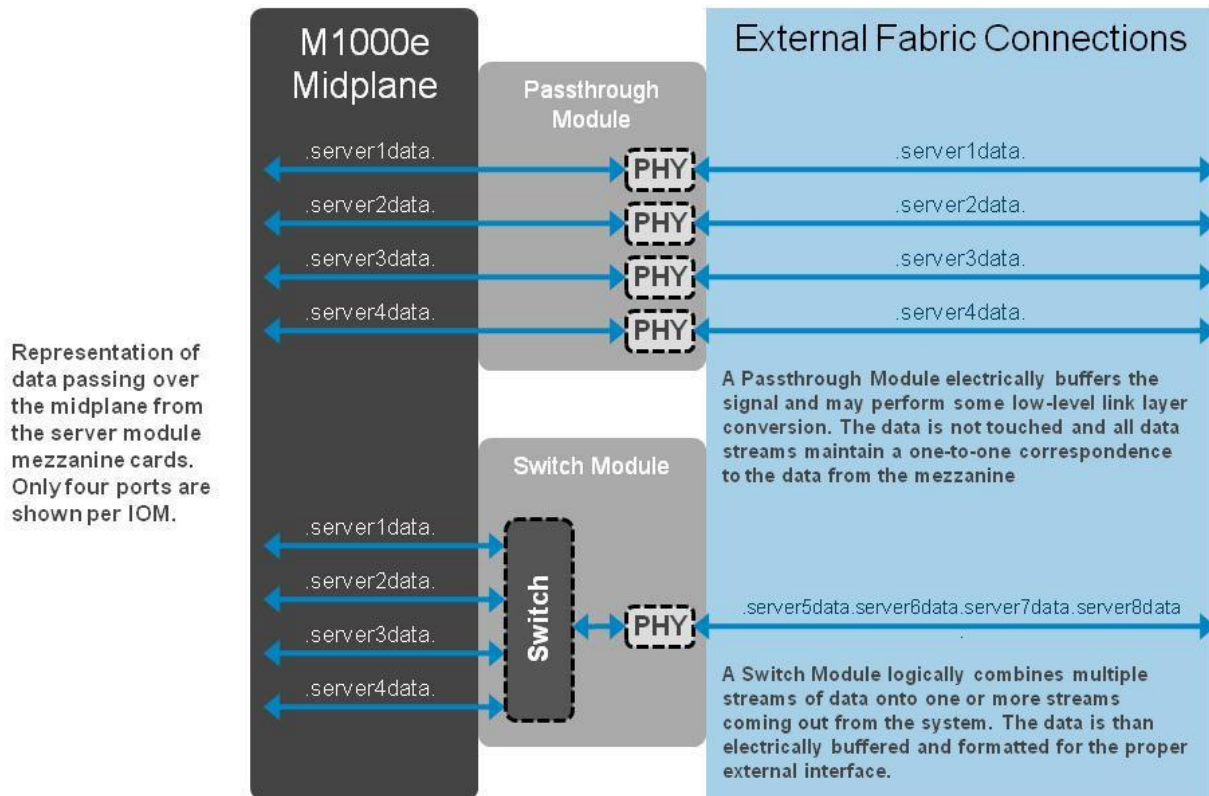


Figure 35. Difference Between Passthroughs and Switch Modules

For more information on the I/O module options, see the [PowerEdge M-Series Blades I/O Guide](#).

9.3 Supported Mezzanine Cards and Switches

Dell supports one mezzanine design standard and one I/O Module design standard for true modular computing.

The currently supported I/O modules include:

- PowerConnect M6220 Switch; GbE + 10GbE uplinks & stacking
- PowerConnect M6348 Switch; 48 1GbE ports + 10GbE uplinks
- PowerConnect M8024 10Gb Ethernet Switch (SFP+, CX4, & 10Gbase-T uplink module options)
- Cisco® Catalyst® 3032 switch; All 1GbE
- Cisco Catalyst 3130g Switch; All 1GbE + stacking
- Cisco Catalyst 3130x Switch; 1GbE+ 10GbE uplinks & stacking
- Cisco 3130g & 3130x switches can be combined in a stack
- 1Gb Ethernet Pass-Through Module
- 10Gb Ethernet Pass-Through Module (SFP+)
- Brocade® 8Gb Fibre Channel Switch

Dell

- 4Gb Fibre Channel Pass-Through
- Mellanox® DDR (20Gb) InfiniBand Switch
- Mellanox QDR (40Gb) InfiniBand Switch

See the [Ethernet I/O Cards](#) page on Dell.com for supported I/O hardware.

9.4 I/O Module Installation

For detailed information on installing the I/O modules in your system, see the I/O Modules section in the *Hardware Owner's Manual* for your specific PowerEdge server.

9.5 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 as was done in the past. By removing the network and storage identity from the server hardware, customers are now able to 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 IO Mezzanine card, upgrade the IO 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

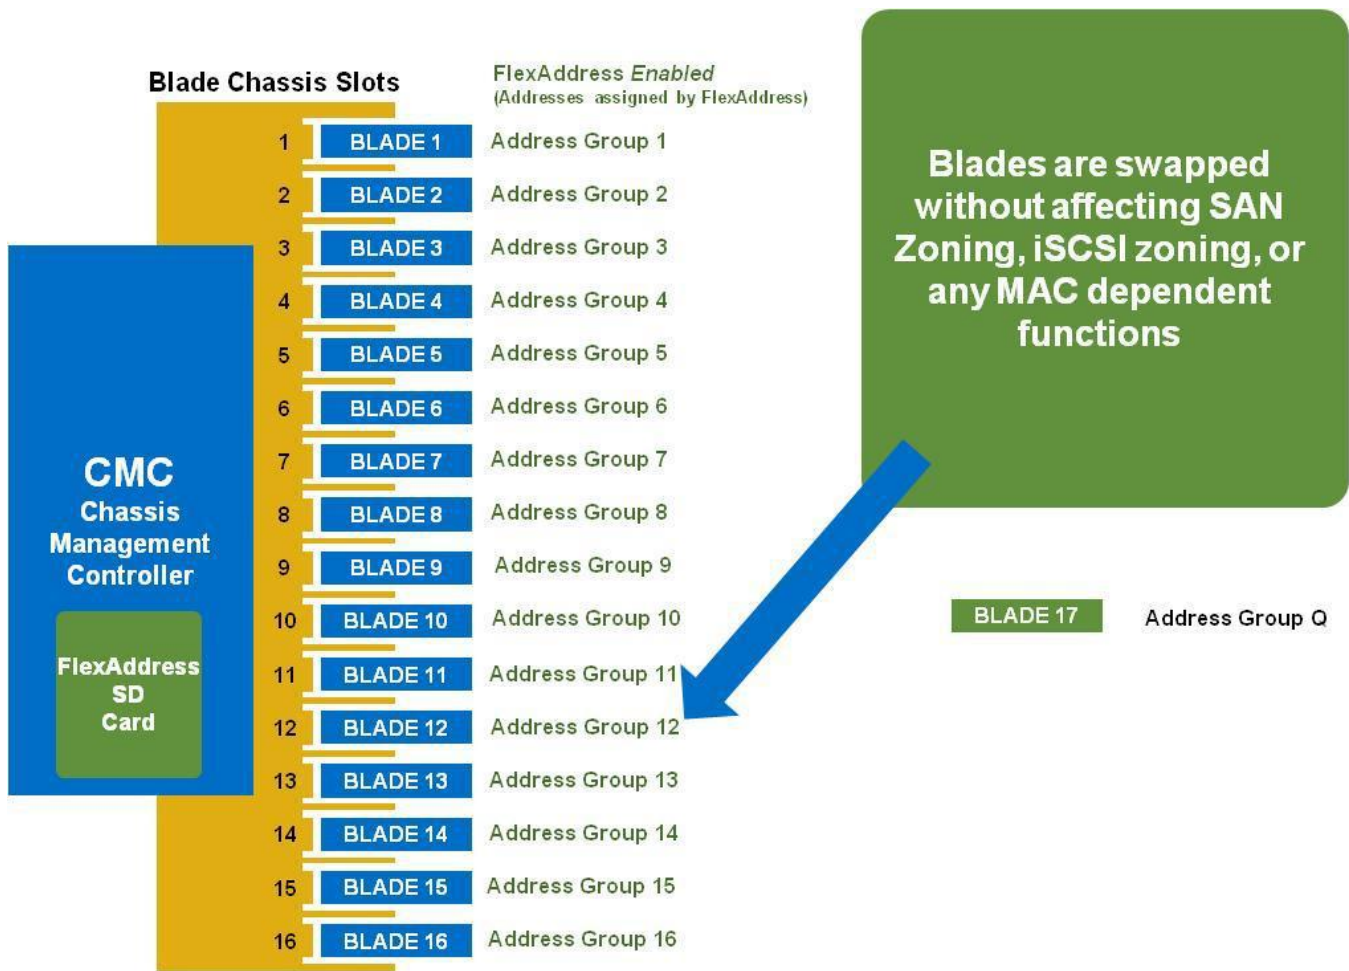


Figure 36. FlexAddress Addresses

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, and thus allows blades to be swapped without affecting SAN Zoning, iSCSI zoning, or any MAC-dependent functions. The write-protected FlexAddress SD card comes provisioned with unique pool of 208 MACs and 64 WWNs. Other types of SD cards inserted into the CMC’s SD card slot are ignored.

FlexAddress can be ordered with a new enclosure or implemented on one already owned by a customer through the purchase of a customer kit. If FlexAddress is purchased with the chassis, it will be installed and active when the system is powered up. In the case of an existing enclosure, FlexAddress requires the addition of one FlexAddress SD card to a CMC and an upgrade to the iDRAC 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. When redundant CMCs are installed, it is not necessary to put such an SD card in both CMCs, since the WWN/MAC addresses are pushed to the chassis Control Panel upon enablement 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 as they already have their WWN/MACs programmed into their controllers. If a replacement of the control panel is required, the SD card will push the WWN/MACs back to it. It is important to note that the chassis Control Panel also stores CMC configuration information, so it is advisable that customers keep a backup of the CMC configuration file.

Dell

The CMC manages the following functions specific to FlexAddress:

- Provides 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, IO 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.

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		

Figure 37. FlexAddress Screen in the CMC

Table 7. FlexAddress Features and Benefits

Features	Benefits
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	Easily replace blades without network management effort
Service or replace 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 & Efficient integration into existing network infrastructure	No need to learn a new management tool Low cost vs switch-based solution
FlexAddress is simple and easy to implement	Simple and quick to deploy
FlexAddress SD card comes with a unique pool of MAC/WWNs and is able to 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-thru modules	Choice is 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 needs to 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 Wake-On-LAN (WOL) function.

Following are the four steps required in order to implement FlexAddress on an M1000e:

- A Chassis Management Controller (CMC) receives a FlexAddress feature card provisioned with a unique pool of: 208 MACs and 64 WWNs.



Figure 38. FlexAddress SD Card

- The FlexAddress card is inserted at factory or to an already installed chassis in a customer location (note: to Enable FlexAddress on an existing chassis, blade BIOS, iDRAC firmware, HBA and Ethernet controller firmware, and CMC firmware must be updated to the latest versions)



Figure 39. SD Slot on bottom of CMC

- Select the slots and fabrics you want FlexAddress enabled on (see figure below). note blades must be powered off and have the latest firmware in order for FlexAddress MAC/WWNs to be deployed to them.

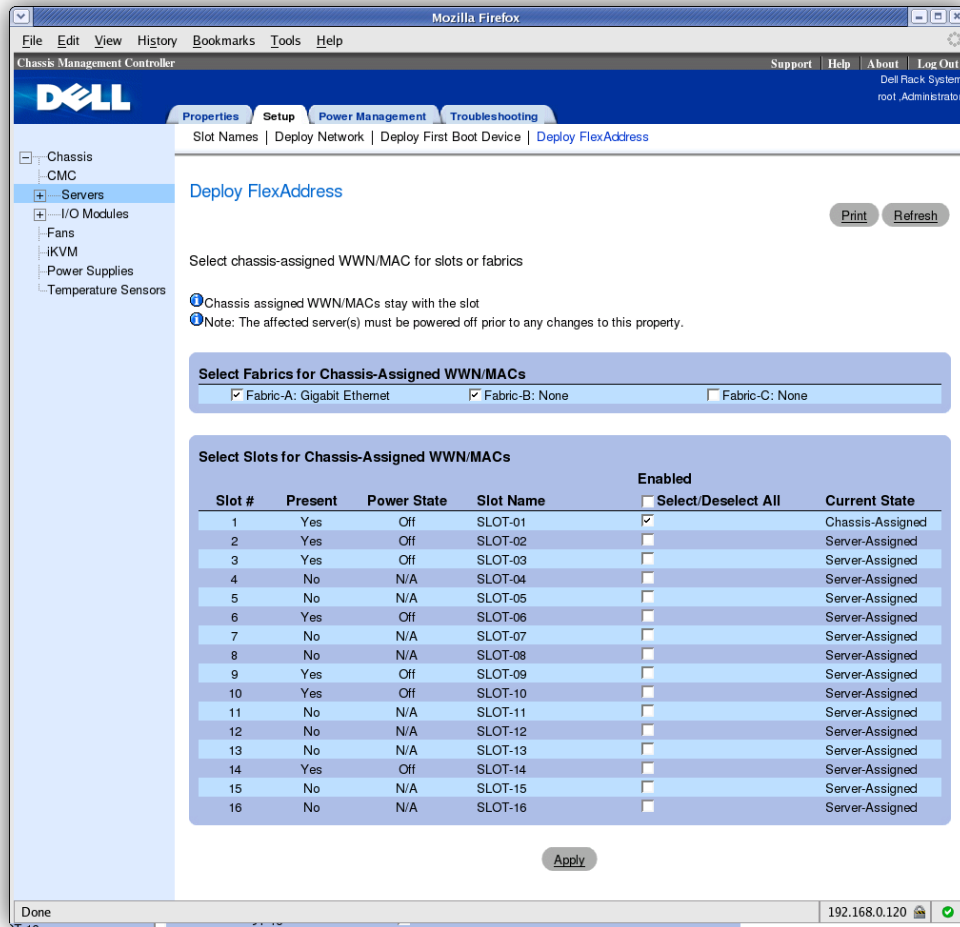


Figure 40. CMC FlexAddress Summary Screen

- FlexAddress MAC/WWNs are now deployed. CMC GUI shows users a summary of Server Assigned (hardware based) and chassis assigned (FlexAddress) MAC/WWNs for the entire chassis or per slot (see Figure Figure 41 below). Green Checkmarks denote which identifier is currently being used.

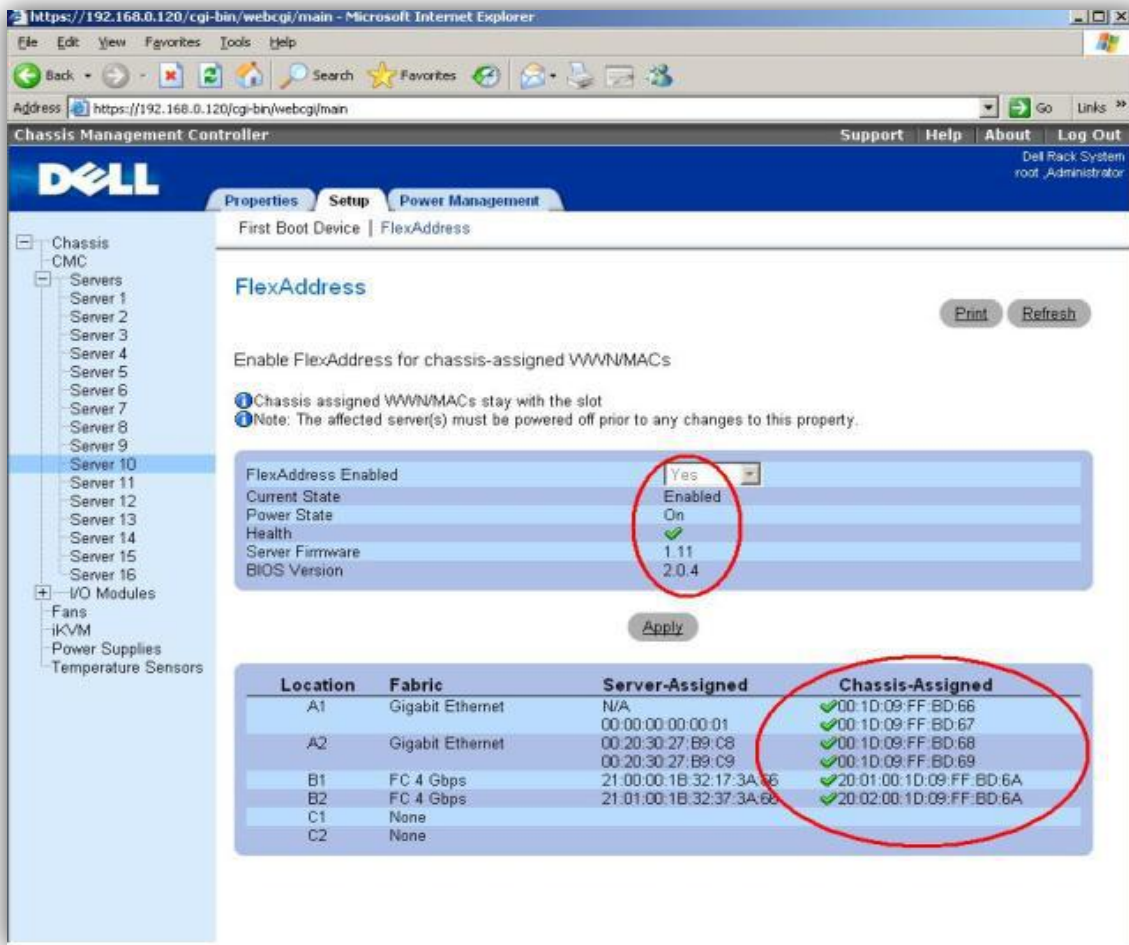


Figure 41. CMC FlexAddress Server Detail Screen

For more information, review the [FlexAddress](#) chapter of the [CMC User's Guide](#).

10 Storage

The M1000e was designed primarily to support external storage over the network, which is the primary requirement for customers seeking maximum density.

MD3000i	AX Series	PS Series	CX Series
 <ul style="list-style-type: none"> • Entry level iSCSI • Consolidation • 16 Servers maximum • Windows, Linux & VMware 	 <ul style="list-style-type: none"> • Entry Fibre / iSCSI • Consolidation • Edge apps • Dell EMC Eco-system 	 <ul style="list-style-type: none"> • Enterprise iSCSI • Consolidation • Ideal for virtualization • Comparable performance to Fibre Channel 	 <ul style="list-style-type: none"> • Enterprise Fibre Channel & iSCSI • Consolidation • Heterogeneous environments • Synchronous replication • Dell EMC Ecosystem

Figure 42. Examples of Major Storage Platforms Supported

The blade servers also support at least two internal hard drives which can be put into RAID if so desired.

For details, see the *Technical Guide* for each of the compatible blade servers offered.

11 Video

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

12 Rack Information

12.1 Overview

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 & 2400 series. 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.

12.2 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.




Figure 43. M1000e RapidRails Static Rails

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 44. M1000e VersaRails Static Rails

The VersaRails static rails are not intended to be mounted in threaded hole racks since 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.

Mounting Interface	Rail Type	Rack Types Supported					Rail Adjustability Range (mm)					
		4-Post			2-Post		Square		Round		Threaded	
		Square	Round	Thread	Flush	Center	Min	Max	Min	Max	Min	Max
RapidRails	Static	✓	✗	✗	✗	✗	712	755	—	—	—	—
VersaRails	Static	✓	✓	✗	✗	✗	706	755	706	755	—	—

The min-max values listed in the table above represent the allowable distance between the front and rear mounting flanges in the rack.

12.3 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 to help manage and secure the potentially large number of cables exiting the back of the system.



Figure 45. M1000e Strain Relief Bar and Cable Enumerator Clip (12 Per Kit)

12.4 Rack View

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 46. M1000e Mounted in the Rack

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 indicated below.




Figure 47. M1000e Strain Relief Bar and Cable Enumerator Clips

More information can be found in the [Rack Installation Guide](#).

13 Virtualization

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



- 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

Figure 48. Examples of Major Virtualization Platforms Supported

The M1000e platform offers many benefits for virtualization:

- Data center Consolidation
 - High Density Form Factor
 - I/O Bandwidth & Switch Port Savings
 - Large Memory Capacity
- Ease Of Mgmt/Deployment
 - Management Options
 - I/O Virtualization
 - Chassis LCD Display
 - Embedded Hypervisor
- Reduce Downtime
 - Persistent Addresses
 - Fully Redundant Power & Cooling
 - Fully Redundant I/O
 - Hot-Swappable Drives
- Power & Cooling Efficiency
 - Super Efficient Power Supplies
 - Optimized Airflow
 - Best-in-class Fan Technology

Many of the I/O modules offered provide significant benefits for use in virtualized environments, such as:





 <p>BROCADE</p> <p>M5424 FC8</p> <p>The world's first end-to-end 8Gb Fibre Channel solution for blades provides more bandwidth to meet your virtualization needs</p>	 <p>CISCO</p> <p>M3130</p> <p>Management simplification, performance, and scalability combine to minimize TCO</p>	  <p>DELL</p> <p>PowerConnect M8024 10GbE</p> <p>10 GbE switch, with 8 external ports for virtualization performance</p> <p>Modular options: SFP+, CX4, 10G Base-T</p>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 49. Examples of I/O modules Recommended for Use in Virtualized Environments

14 Systems Management

14.1 Overview

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

- Dual Redundant Chassis Management Controllers (CMC)
 - 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 like 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 admin per blade
 - Control Panel on front of M1000e for “crash cart” access
- Front LCD
 - Designed for deployment and local status reporting

Onboard graphics and keyboard/mouse USB connect to an optional system level Integrated KVM (iKVM) module for local KVM access. Full USB access is available through the server module front panel. In Figure 50 below, OSCAR (On Screen Configuration and Activity Reporting) is the graphic user interface for the 4161DS or 2161DS-2 console switch used to share a single keyboard/mouse/LCD in a rack with all the servers in the rack

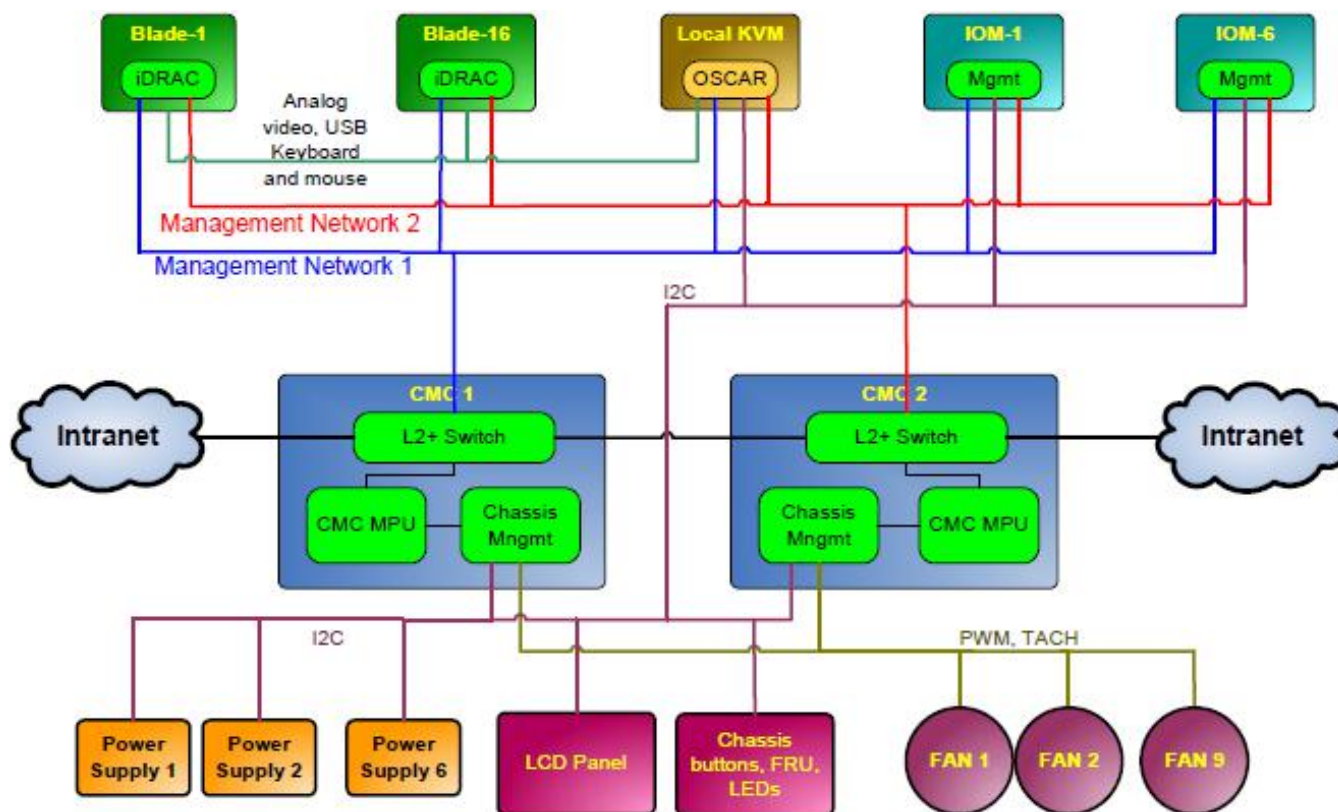


Figure 50. System Management Architecture Diagram

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.

14.2 Server Management

The server module base management solution includes additional features for efficient deployment and management of servers in a modular server form factor. The base circuit, which integrates the baseboard management controller (BMC) function with hardware support for Virtual KVM (vKVM) and Virtual Media (vMedia), is the integrated Dell Remote Access Controller (iDRAC). iDRAC has two Ethernet connections, one for each CMC, providing system management interface redundancy.

Highlights of the iDRAC solution include the following:

- Dedicated management interface for high-performance management functions
- vMedia
- vKVM
- IPMI 2.0 Out-of-Band management
- Serial-over-LAN redirection
- Systems Management Architecture for Server Hardware (SMASH) Command Line Protocol (CLP)
- Blade status and inventory
- Active power management

Dell

- Integration with Microsoft® Active Directory®
- Security, Local and Active Directory

Traditional IPMI-based BMC features like hardware monitoring and power control are supported.

The Lifecycle controller on 11th generation servers offers additional features including:

- Unified Server Configurator (USC): Consolidated interface for OS install, hardware configuration, updates, and diagnostics
 - Reduces task time and speed of deployment
 - Eliminates media for OS drivers during OS install and update
 - Eliminates multiple control ROM options for hardware configuration
 - Eases firmware updates with roll-back capability
- Embedded Diagnostics
- Parts Replacement: Automatically updates RAID/NIC firmware to previous levels after parts replacement.
- Persistent Life Cycle Log
- Out-of-Band Configuration/Update: With the introduction of Lifecycle Controller (LCC) 1.3, customers and console vendors can initiate and schedule an upgrade to device firmware and pre-OS software out-of-band. This is useful during “bare metal” deployment scenarios or change management in operating systems where firmware updates were not possible in the past (i.e., virtualization).
 - Updates are staged on the Lifecycle Controller and applied immediately or during a scheduled maintenance window.
 - Lifecycle controller 1.3 supports updates with BIOS, diagnostics, driver pack, USC, RAID controller firmware, iDRAC6 firmware, and NIC firmware.

More information on the [iDRAC](#) and [Lifecycle Controller](#) can be found on support.dell.com.

14.3 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 rear 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
 - 1:Many iDRAC configuration
 - 1:Many Blade Boot Device Selection
 - 1:Many vMedia File share
 - Customized Slot Naming
 - IO Module Configuration and Launch
 - WWN/MAC Display & Persistence w/ FlexAddress; Manages FlexAddress Ports
 - Support for Network Time Protocol (NTP)
- Monitoring and Troubleshooting
 - User interface entry point (web, telnet, SSH, serial)

Dell

- 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 & Event Logs
 - Email & SNMP alerting
 - Support for Remote Syslog
 - Blade Events displayed in CMC
- Consolidated Chassis/Blade/IO Inventory
- “Virtual Server Reseat” simulates blade removal/insertion
- Remotely blink LEDs to Identify Components
- Updating
 - Reporting of Firmware versions
 - 1:Many iDRAC Firmware update
 - Consolidated CMC and iKVM F/W update
 - 1:Many update of drivers & Firmware via Remote File Share (w/ Repository Manager)
 - 2 x 10/100/1000Mb Ethernet ports + 1 serial port
- Real Time Power/Thermal Monitoring and Management
 - Consolidated Chassis/Blade Power Reporting
 - Power budget management and allocation
 - Real Time System AC Power Consumption with reset-able peak and minimum values
 - System level power limiting and slot based power prioritization
 - Manages Dynamic Power Engagement functionality
 - Manages fan speed control
 - Power sequencing of modules in conjunction with the defined chassis power states
- Separate management network
 - Configuration of the embedded management switch, which facilitates external access to manageable modules
 - Provides connection from management network to iDRAC on each of the blades and the management interfaces on the integrated I/O Modules
 - 2nd Ethernet port supports daisy chaining of CMCs for improved cable management
- Security
 - Local Authentication &/or AD Integration
 - OpenLDAP coming in CMC 3.0
 - Supports multiple levels of user roles and permissions for control of chassis, IO, &/or server blades, including Microsoft Active Directory
 - IPv6 Support
 - VLAN tagging for iDRAC, CMC, and IOMs
 - Two Factor Authentication w/ 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


Dell

- SNMP

The Integrated Dell Remote Access Controller (iDRAC) on each server module is connected to the CMC via 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 51. Chassis Management Controller



- | | | | |
|---|-----------------------------------------------|----|------------------------------------------------------------------------------|
| 1 | Ethernet connector Gb1 | 2 | Ethernet connector STK (used for daisy-chaining CMCs in separate enclosures) |
| 3 | link indicator (2) | 4 | activity indicator (2) |
| 5 | DB-9 serial connector for local configuration | 6 | optional secondary CMC (CMC 2) |
| 7 | primary CMC (CMC 1) | 8 | amber fault indicator |
| 9 | blue status/identification indicator | 10 | power indicator |

Figure 52. CMC Module Features


14.4 Integrated Keyboard and Mouse Controller (iKVM)



Figure 53. M1000e iKVM

The modular enclosure supports one optional Integrated KVM (iKVM) module. This module occupies a single slot accessible through the rear 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 via 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 the CMC Command Line Interface via the KVM switch and allowing text-based deployment wizards on VGA monitors. iKVM firmware is updated through the CMC.



- | | | | |
|---|--------------------------------------|---|-------------------------------------------|
| 1 | identification indicator | 2 | status indicator |
| 3 | ACI port for tiering connection only | 4 | USB connectors (2) for keyboard and mouse |
| 5 | video connector | | |

Figure 54. Rear iKVM interface Panel

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 55. 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

Dell

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\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, customers need to connect to the card using the USB adapter.

More information on the iKVM can be found in the [iKVM training](#) material on [us.training.dell.com](#), in the [iKVM Module](#) section of the [Dell PowerEdge Modular Systems Hardware Owner's Manual](#), in the [CMC User Guide](#), and at [dell.avocent.com](#).

15 Peripherals

Common peripherals for the M1000e include:

- An external USB DVD-ROM Drive is often used for local installation of OS or other software.
- A Dell 1U rack console which enables customers to mount a system administrator's control station directly into a Dell rack without sacrificing rack space needed for servers and other peripherals. It features:
 - 17" 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 for racks, which provides a 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 your servers
 - Support for up to 5000 VA (3750 watts)
- Power distribution units (PDUs): use the [Dell Energy Smart Solution Advisor](#) (ESSA) to see what a given chassis configuration will require.
 - Single phase needs one PDU per chassis
 - Use 30A for a medium to lightly loaded chassis
 - Use 60A for a heavily loaded
 - For 3 phase:
 - 30A 3 phase for a heavily loaded single chassis
 - 50 or 60A 3 phase for multiple chassis.
 - Generally customers run 2 x 3 phase circuits to the rack via a PDU, then the PDU breaks out single phases to each PSU.

16 Packaging Options



Figure 56. Enclosure After Unpacking

The M1000e comes from the factory on a pallet with components installed. The components must be taken out before it is installed in a rack. There are cardboard slings under the chassis to enable two people to lift it.

Dell

Appendix A. Regulatory Certifications

Please see the external Product Safety, EMC, and Environmental Datasheets on dell.com at:
http://www.dell.com/regulatory_compliance_datasheets.

Dell

Appendix B. Status Messages

C.1 LCD Status Messages

See the [M1000e Configuration Guide](#) and [CMC Administrator Reference Guide](#) for details.

C.2 System Status Messages

See the [CMC Administrator Reference Guide](#) for details.

Dell

Appendix C. Additional Information

Videos highlighting the major M1000e features are available on http://www.dell.com/html/us/products/pedge/poweredge_mseries_blade_videos/poweredge.html

The [PowerEdge M1000e Configuration Guide](#), [Hardware Owner's Manual](#) and [CMC Administrator Reference Guide](#) each contain a wealth of additional information about the PowerEdge M1000e's capabilities.

The [blade training material on dtt.us.dell.com](http://blade.training.material.on.dtt.us.dell.com) is also a useful reference.