PowerEdge M1000e

Technical Guide

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

Revision 1 November 2011

Table of Contents

Tables

Figures

1 Product Comparison

The Dell™ PowerEdge™ M1000e modular server enclosure 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/M610/M610x/M710/M710HD/M910/M915
Form Factor	7U	10U
Max No. of Blades	10	16 half-height or 8 full-height
I/O Module Bay	4	6
Fabric Types Supported	1 x Dual GbE1 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	Up to 6 x 2700W PSUs in either non-redundant, PSU redundant or grid redundant profiles
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.

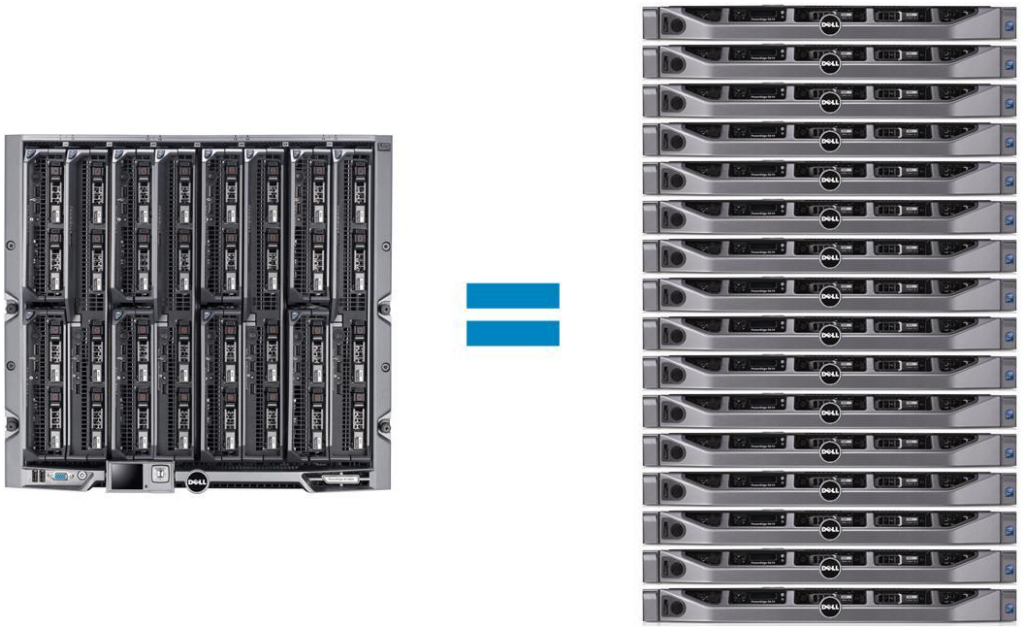

Figure 1. Server Density Comparison

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 Table 2 (darker blue shading indicates increased density).

Table 2. Rack vs. Blade Server Rack-Level Specification Comparison[1]

	R410	R510	R610	R710	R810	R815	R910	M610	M610x	M710	M710HD	M910
Form Factor	Rack[2]	Rack	Rack	Rack	Rack	Rack	Rack	1/2 Blade	Full Blade	Full Blade	1/2 Blade	Full Blade
Manufacturer	Intel	Intel	Intel	Intel	Intel	AMD	Intel	Intel	Intel	Intel	Intel	Intel
Sockets	2	2	2	2	4	4	4	2	2	2	2	4
Max Cores per 42U Rack	504	252	504	252	672	1,008	320	768	384	384	768	1,024
Max RAM per rack, in TB	5	3	8	4	11	5	10	12	6	6	12	16
Max 1GB Ethernet Ports per 42U Rack	252	378	504	420	588	588	440	640	320	576	640	640
Max 10GbE, DDR IB, or FC network ports per rack	84	168	168	84	252	252	220	256	128	256	256	256
Drives per 42U rack	168	252	252	168	126	126	160	128	64	128	128	64
Max 7.2k or 10k rpm internal storage per rack	336	504	151	252	76	76	96	77	38	77	77	38
Max 15k rpm internal storage per 42U rack	101	151	37	76	18	18	23	19	9	19	19	9
Max SSD internal storage per 42U rack	17	25	25	17	13	13	16	13	6	13	13	6

[1] This rack-level physical capacity specification summary does not factor in power and cooling.
[2] 42U is the most common rack size.

2 Key Technologies

2.1 Overview

The Dell™ 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, and 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 and 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:

- A high-speed passive midplane that connects the server modules in the front and to 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 (with 4x QDR InfiniBand) with future support for higher-bandwidth I/O devices when those technologies become available; this provides high-speed server module connectivity to the network and to 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/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 six hot-pluggable, redundant power supplies and nine hot-pluggable, N+1 redundant fan modules
- System front control panel with LCD panel and two USB keyboard/mouse connections and one video crash-cart connection

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 2000 mm depth racks.

3.2 Product Features Summary

Table 3. Product Features Summary

Feature	Parameter
Chassis Size	10U (height) 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	Up to 6
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 (each configuration 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)	3+3 2+2 1+1
DC Redundancy (each configuration has one extra power supply that comes online if one of the existing power supplies fails)	1+1 2+1 3+1 4+1 5+1

4 Mechanical

4.1 Chassis Description

The Dell™ PowerEdge™ 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 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 the System Control Panel Features section in the About Your System chapter of the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on Support.Dell.com/Manuals. M1000e supports up to sixteen half-height or eight 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

Table 4. Dimensions

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

4.3 Front View and Features

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 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 the System Control Panel Features section in the About Your System chapter of the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on Support.Dell.com/Manuals.

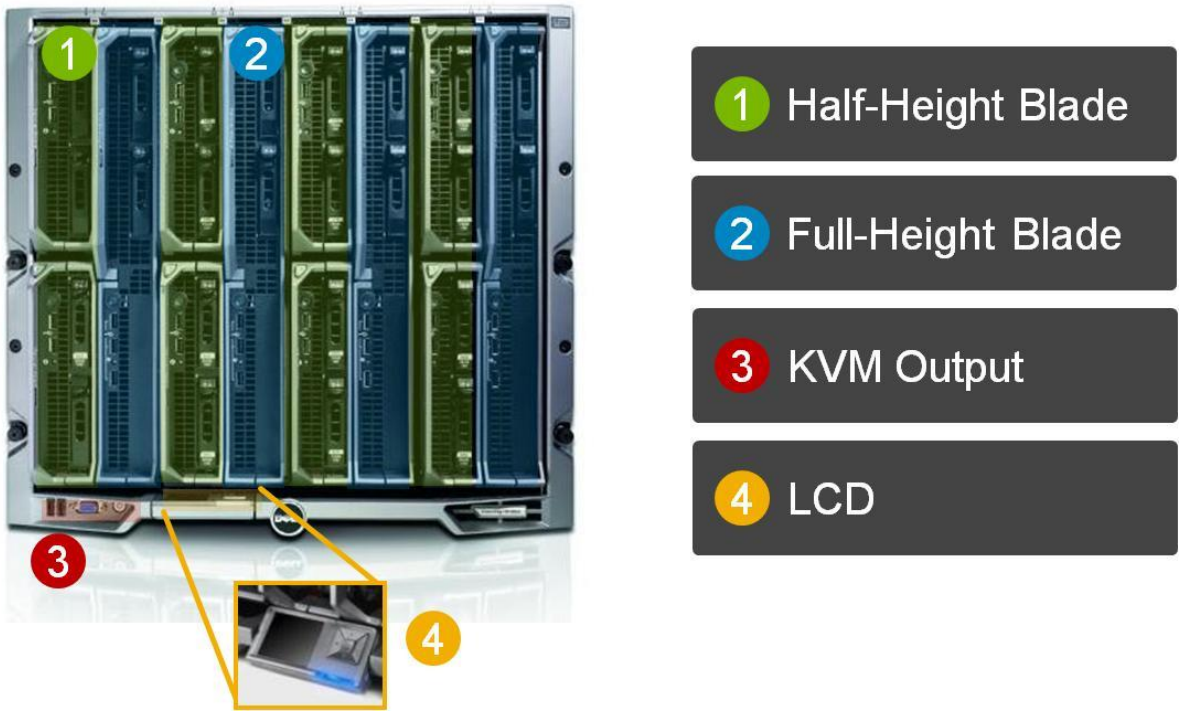

Figure 2. M1000e Front View

The M1000e enclosure supports half-height or full-height server modules, as well as other form factors, including double-width modules.

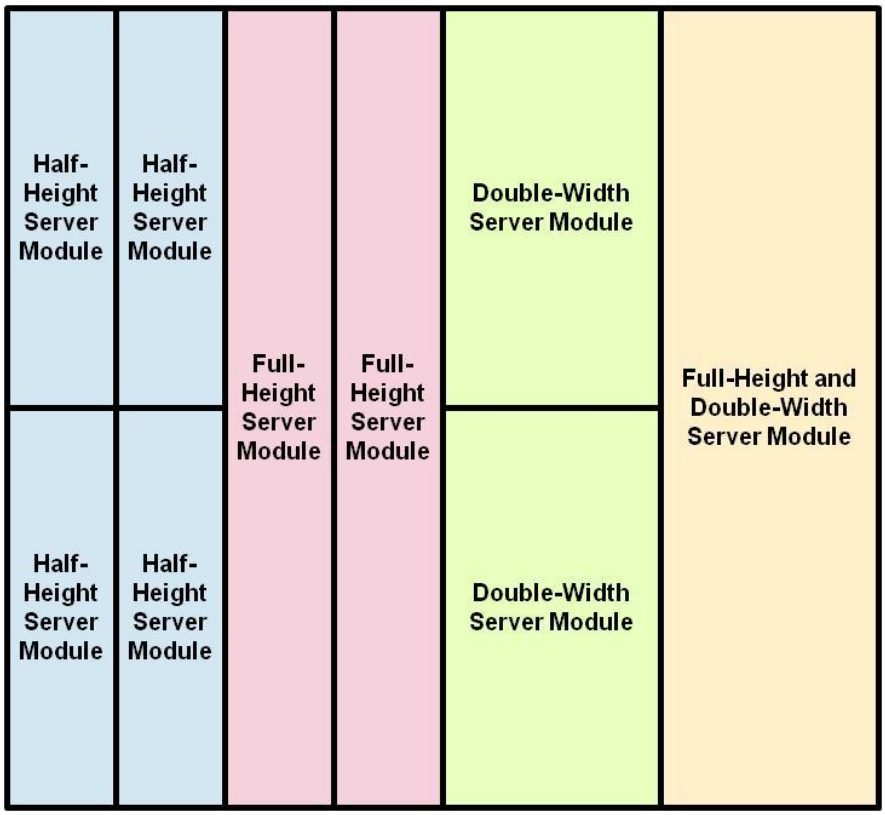

Figure 3. Possible Server Module Sizes, Front View

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.

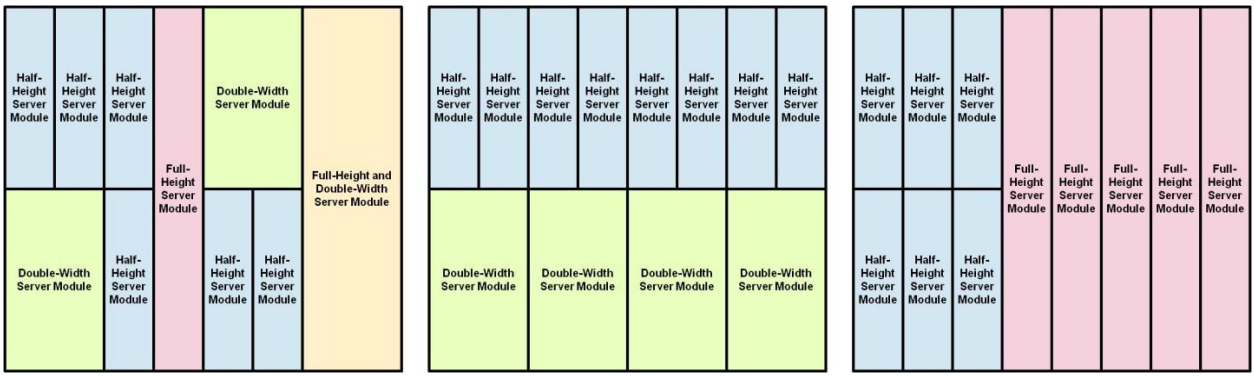

Figure 4. Example Server Module Configurations

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.

4.4 Back View and Features

Figure 5. Back View

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. For more information, see the Back Panel Features section in the About Your System chapter of the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on Support.Dell.com/Manuals.

4.5 Power Supply Indicators

For information on M1000e power supply indicators, see the Power Supply Indicator section in the About Your System chapter of the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on Support.Dell.com/Manuals.

4.6 Rails and Cable Management

RapidRails™ Static Rails for square-hole racks support the following:

- Toolless installation in 19" 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 712-755 mm
- Minimum rail depth of 703 mm

VersaRails™ Static Rails for square- or round-hole racks support the following:

- Tooled installation in 19" 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 706-755 mm
- Round-hole rack adjustment range of 706-755 mm
- Minimum rail depth of 703 mm

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 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	—	—	—
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

For information on cabling for the M1000e modular system, see the Routing and Managing Cables section in the *Rack Installation Guide* on Support.Dell.com/Manuals.

RapidRails™ Static Rails for square-hole racks support toolless 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. 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. For information on the RapidRails and VersaRails rack kits, see the General Installation Instructions section in the *Rack Installation Guide* on Support.Dell.com/Manuals.

See Section 12 for more rack information for the M1000e.

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 rack information for the M1000e.

4.8 Rack View

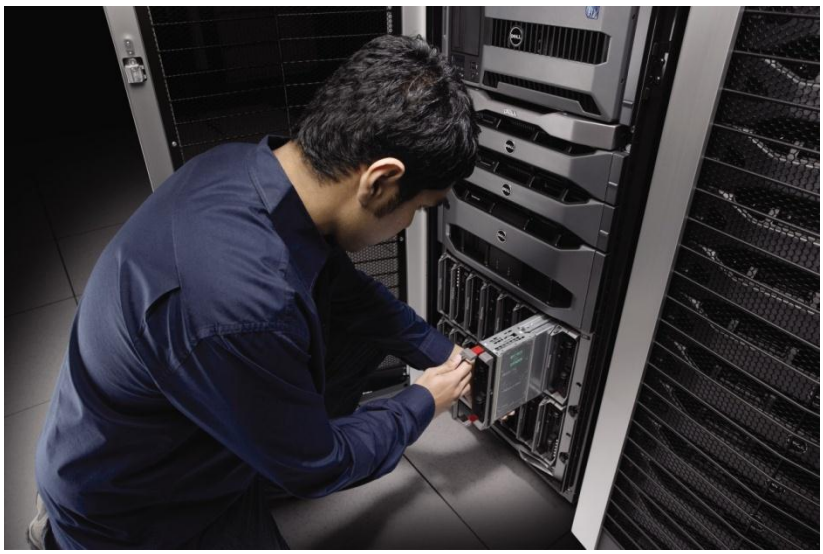

Figure 6. M1000e in a Rack

4.9 Fans

Figure 7. Back View of Chassis Showing Fans

The PowerEdge M1000e chassis comes standard with nine 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 the following:

- 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 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 8. Blades, Blanks, and One Open Slot Needing to be Filled

Figure 9. Power Supply, Power Supply Blanks, and One Open Slot Needing to be Filled

Figure 10. I/O Module and One Open Slot Needing to be Filled

Figure 11. Installed CMC, I/O Module, and Power Supply Blanks

Figure 12. Installed iKVM Blank

Figure 13. 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 and 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

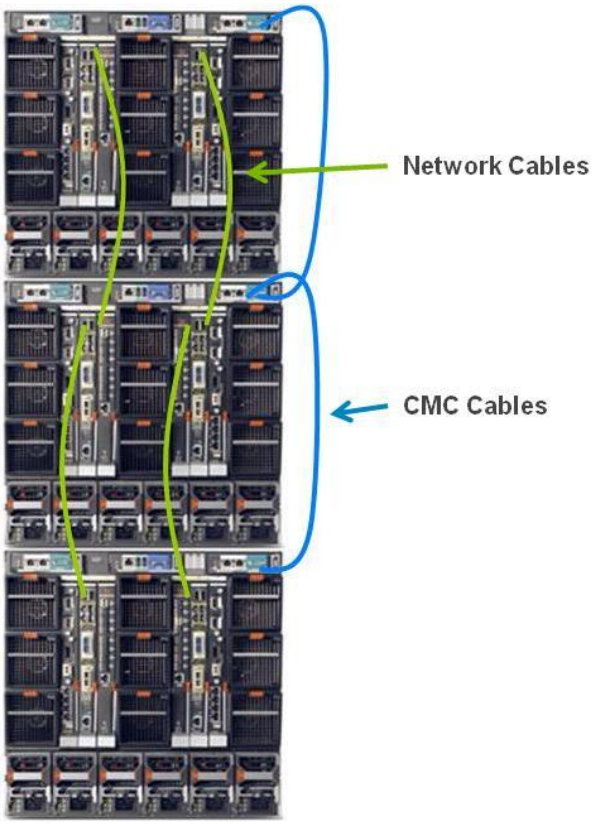

Figure 14. Simplified Cabling

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 15. M1000e LCD Panel in Recessed Position

Figure 16. 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 using the panel, including their type, user-defined name, configurations, service tag numbers, and IP address 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 17 shows some of the capabilities of the LCD control panel.

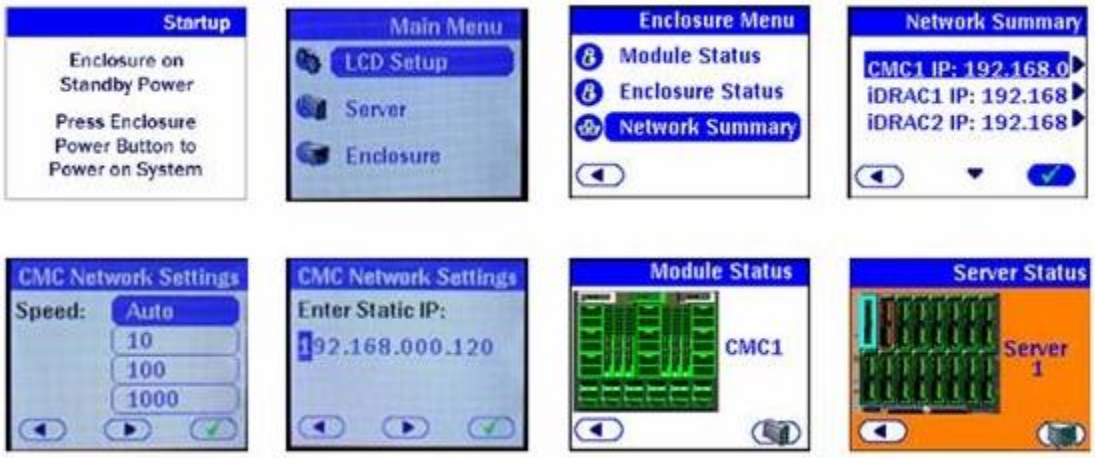

Figure 17. LCD Panel Capabilities

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 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

See the *Dell PowerEdge M1000e Configuration Guide* and the *Dell Chassis Management Controller Firmware Version x.x Administrator Reference Guide* or the *Dell Chassis Management Controller Firmware Version 3.1 User* 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 administrator per blade or one administrator per multiple blades
- Grant permissions to some blades but not to others
- Customize administrative access for CMC, iDRAC, and I/O

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 PowerEdge™ 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 back of the chassis.

Figure 18. Power Supplies in M1000e

Figure 19. M1000e Power Supply Back View

The 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 Dynamic Power Supply Engagement (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, 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 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.

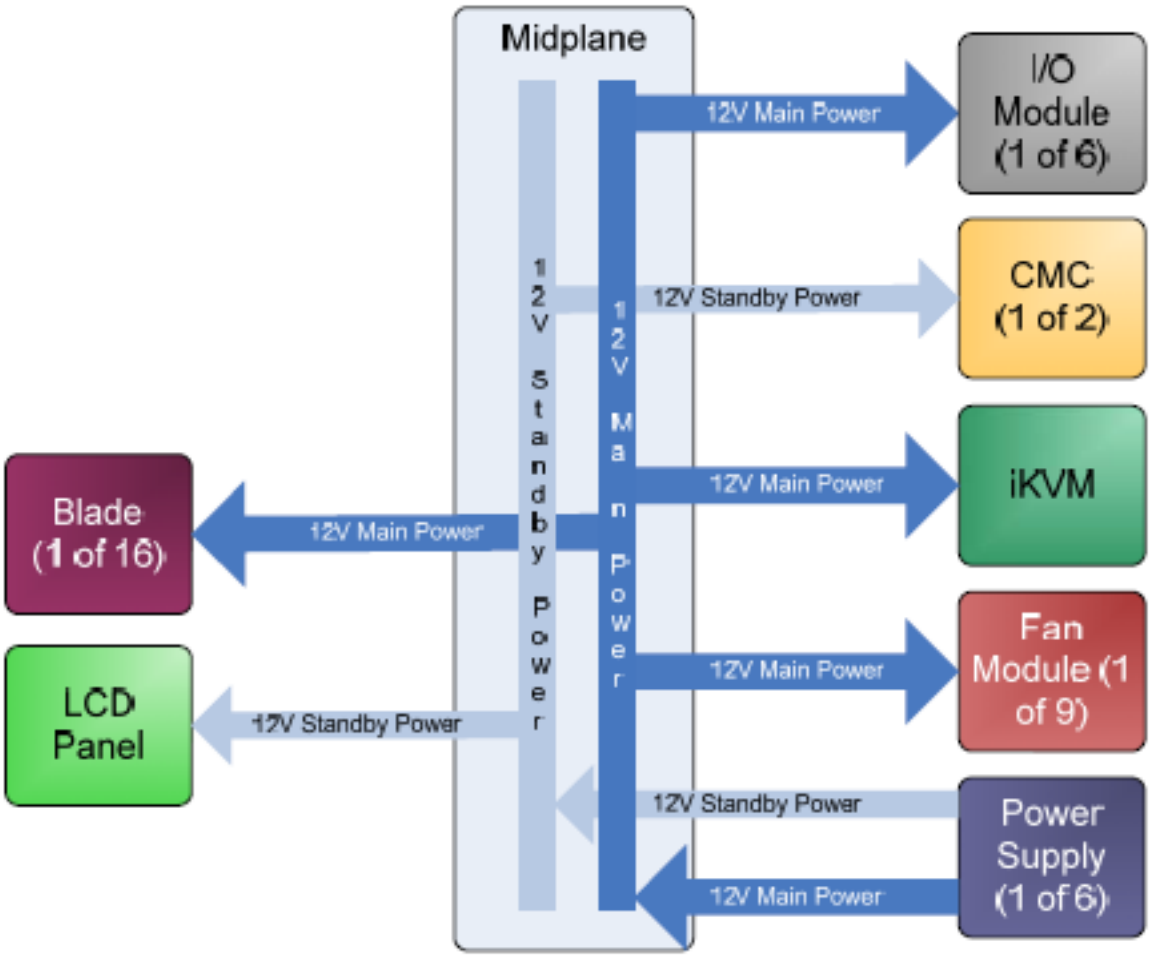

Figure 20. Power Architecture

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

5.1.1 2360 Watt Power Supplies

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.

5.1.2 2700 Watt Power Supplies

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.

5.2 Power Redundancy Options

The six power supplies allow for two different redundancy options: Grid Redundancy and Power Supply Redundancy.

For information on the Grid Redundancy option, see the Power Management chapter in the *Dell Chassis Management Controller Firmware Version x.x User* Guide on Support.Dell.com/Manuals.

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.

5.2.1 Supported Redundancy Options in CMC 3.0

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 2+0 3+0 4+0 5+0 6+0	1+0 2+0 3+0 4+0	1+0 2+0 3+0 4+0 5+0
AC Redundancy	1+1 2+2 3+3	1+1 2+2 3+3	1+1 2+2 3+3
Power Supply Redundancy	1+1 2+1 3+1 4+1 5+1	1+1 2+1 3+1 4+1	1+1 2+1 3+1 4+1 5+1

5.2.2 Power Tables for Redundancy Options

The following tables 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	9922	4920	8888
5+1	9922	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	9922	4920	8888
5+0	9922	6150	9922
6+0	9922	7380	9922

5.2.3 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 will provide 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.

5.2.3.1 Allowable Power Supply Combinations

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

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

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

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

5.2.4 Chassis Behavior When Powering Up Power Supplies

For information on M1000e chassis behavior when powering up the power supplies, see the Power Management chapter in the *Dell Chassis Management Controller Firmware Version x.x User's Guide* on Support.Dell.com/Manuals.

5.2.5 CMC Revisions and Power Supply Compatibility

The following table 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	3.1
2360W	✓	✓	✓	✓	✓
1350W	X	X	X	✓	✓
27000W	X	X	X	✓	✓

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.

5.2.5.1 Behavior of 2700W Power Supply

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

5.2.5.2 Behavior of 1350W Power Supply

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).

5.2.5.3 Behavior of 2700W + 2360W Power Supplies

Mixing 2700W and 2360W power supplies is allowed, but the CMC must be upgraded to 3.0 for the 2700W power supply to contribute power to the chassis. When mixing different capacities, the CMC adjusts the redundancy algorithm to accommodate the different capacities.

5.2.5.3.1 Grid Redundancy

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.

5.2.5.3.2 Power Supply Redundancy

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

5.2.5.4 Behavior of 1350W + 2360W/2700W Power Supplies

Mixing 1350W (110VAC) and 2360W/2700W (220VAC) power supplies is not allowed.

- 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.

Upon 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.

5.2.6 Recommended Number of Servers

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

5.2.6.1 Half-height Blade Servers

Table 11 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 11. Half-Height Servers Per Power Supply Configuration

<table>
<tr><th>Power Supply Configuration</th><th>Number of Servers
220–240V power
2700W PSUs</th><th>Number of Servers
220–240V power
2360W PSUs</th><th>Number of Servers
100–120V power
2700W/1350W PSUs[1]</th></tr>
<tr><td>2+1
2+2</td><td>16 light,
13 medium, or
10 heavy</td><td>16 light,
11 medium, or
8 heavy</td><td>9 light,
4 medium, or
3 heavy</td></tr>
<tr><td>3+0
3+1
3+3</td><td rowspan="3">Any server combination</td><td>16 light,
16 medium, or
15 heavy</td><td>16 light,
8 medium, or
6 heavy</td></tr>
<tr><td>4+1
4+2</td><td rowspan="2">Any server combination</td><td>16 Light,
13 Medium, or
10 Heavy</td></tr>
<tr><td>5+1</td><td>16 light,
16 medium, or
13 heavy</td></tr>
</table>

[1]With 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 11, 1–3 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 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 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.

5.2.6.2 Full-height Blade Servers

0 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 12. Full-Height Servers Per Power Supply Configuration

Power Supply Configuration	Number of Servers 220–240V power 2700W PSUs	Number of Servers 220–240V power 2360W PSUs	Number of Servers 100–120V power 2700W/1350W PSUs[1]
2+1 **2+2**	8 light, 7 medium, or 5 heavy	8 light, 5 medium, or 4 heavy	3 light, 2 medium, or 1 heavy
3+0 **3+1** **3+3**	Any server combination	8 light, 8 medium, or 7 heavy	6 light, 4 medium, or 3 heavy
4+1 **4+2**		Any server combination	8 light, 6 medium, or 4 heavy
5+1			8 light, 8 medium, or 6 heavy

[1]With 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 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 0, 1–3 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 3.02 enables automatic retries if blades attempt simultaneous power-ups and do not power on.

Dell

Dell recommends that all customers running PowerEdge M910 blade servers should install the modular iDRAC 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 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.

Also, be aware that Dell recommends that all customers running M910 should install the modular iDRAC firmware 3.02 upgrade. It contains several enhancements:

- Improved power budgeting to lower the amount of power budget M910 requests from the CMC, enabling more servers to run on fewer power supplies
- Introduction of functionality to enable iDRAC to retry powering up of blade servers automatically in the event some servers do not power up initially due to the pre-boot power inventory

5.3 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 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 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 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 processer 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 will start throttling down the 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.

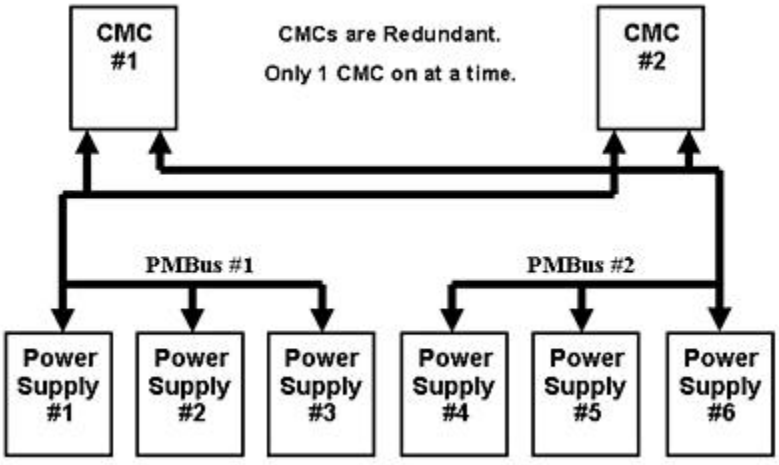

Figure 21. PMBus Communication Channels

Each PSU contains a FRU and a microcontroller. The FRU stores and transfers information about the PSU across the PMBus (I2C bus):

- Vendor part number
- Revision serial number
- Power capability

The microcontroller performs the following functions:

- Control for the main power regulator
- Reports PSU status across the PMBus
  - Input/Output Voltage/Current
  - AC power consumption
  - DC power consumption
  - Error status
  - Individual Status Signals from the PSU 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 in eleventh generation servers. 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 Power Management chapter in the *Dell Chassis Management Controller Firmware Version x.x User's Guide* on Support.Dell.com/Manuals.

5.4 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).
- Dynamic Power Supply Engagement, 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)

5.5 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 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.

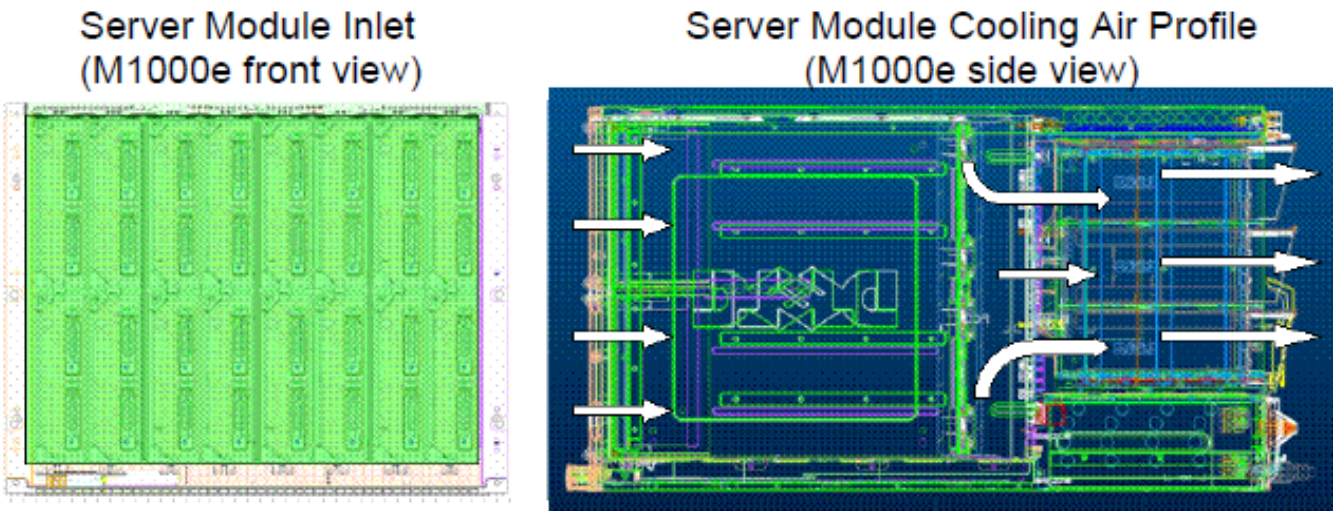

Figure 22. Server Cooling Air Profile

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.

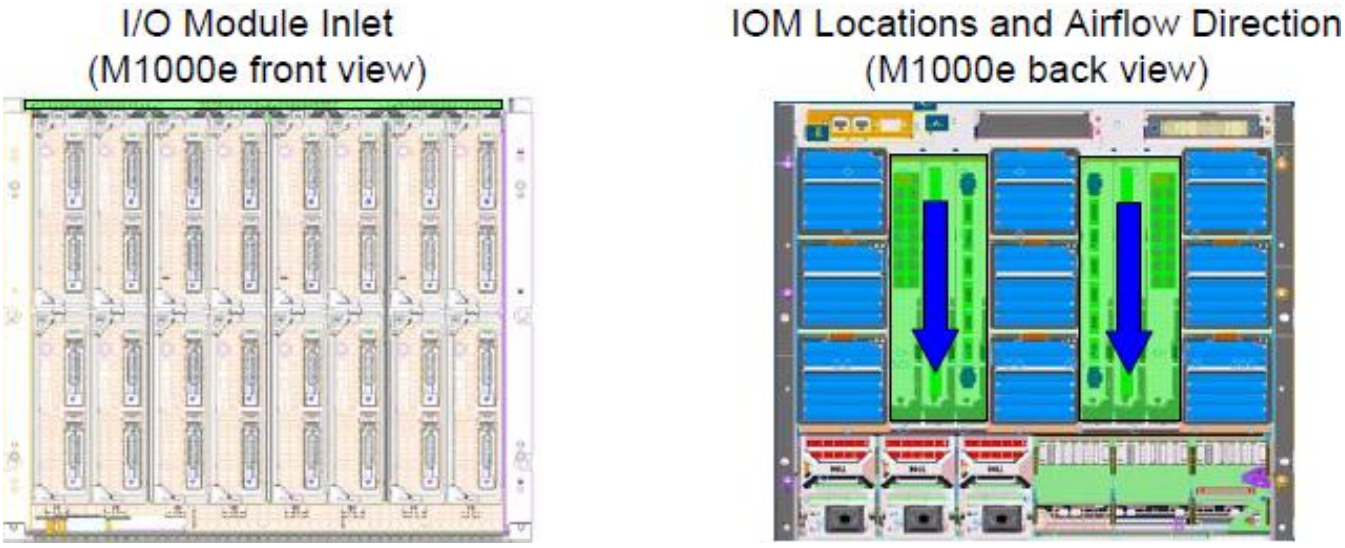

Figure 23. I/O Module Inlet and IOM Locations

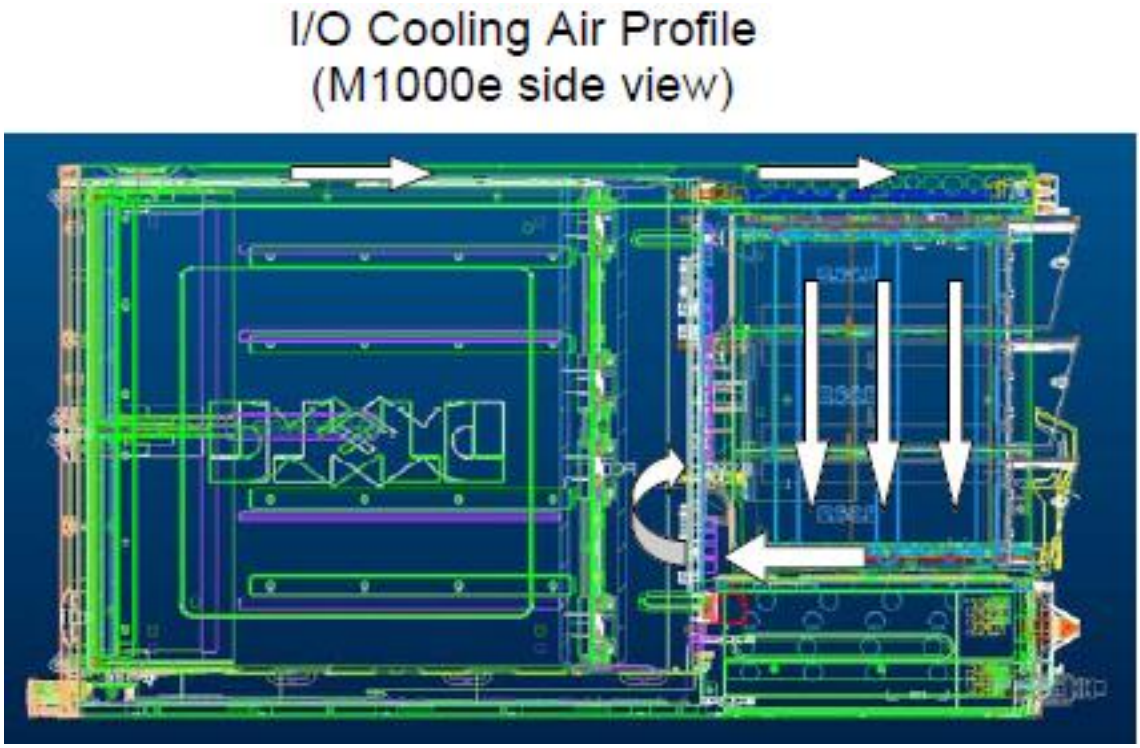

Figure 24. I/O Cooling Air Profile

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.

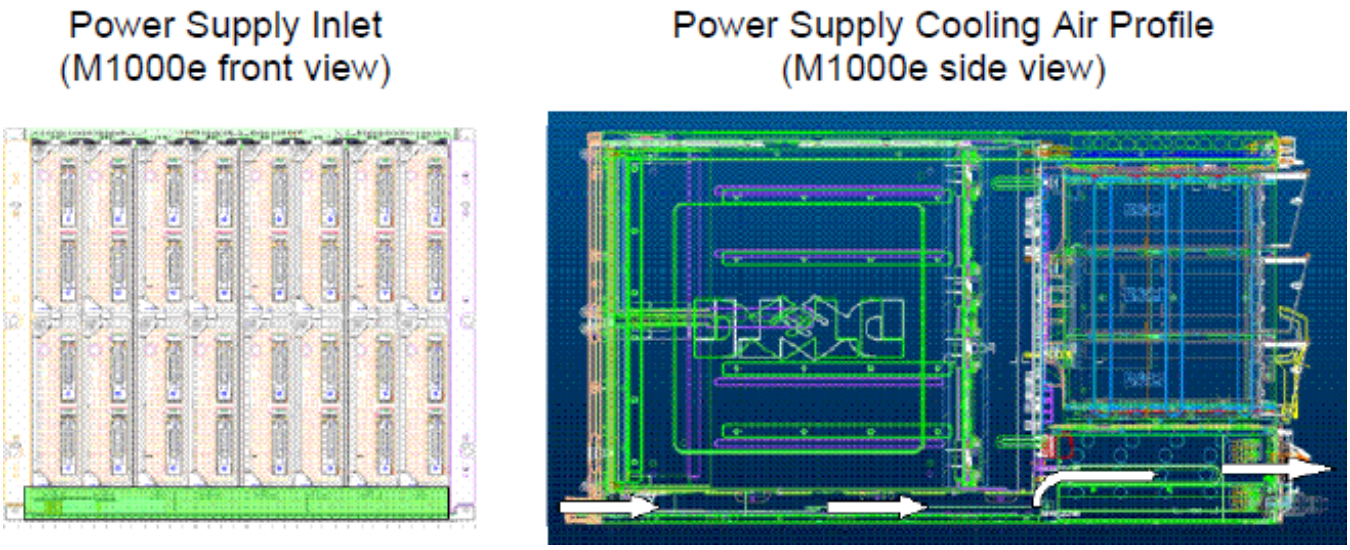

Figure 25. Power Supply Inlet and Cooling Air Profile

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 the figure above. This ensures that the power supplies receive ambient temperature air.

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.

5.6 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 *Getting Started Guide* on Support.Dell.com/Manuals.

5.7 Power Consumption

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

5.8 Maximum Input Amps

See Power Distribution Systems for the Dell M1000e Modular Server Enclosure – Selection and Installation.

5.9 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 iDRACs 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.

5.10 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 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

See the *Dell PowerEdge Mxxx Technical Guide* on Dell.com 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 Dell™ PowerEdge™ 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 26. 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.

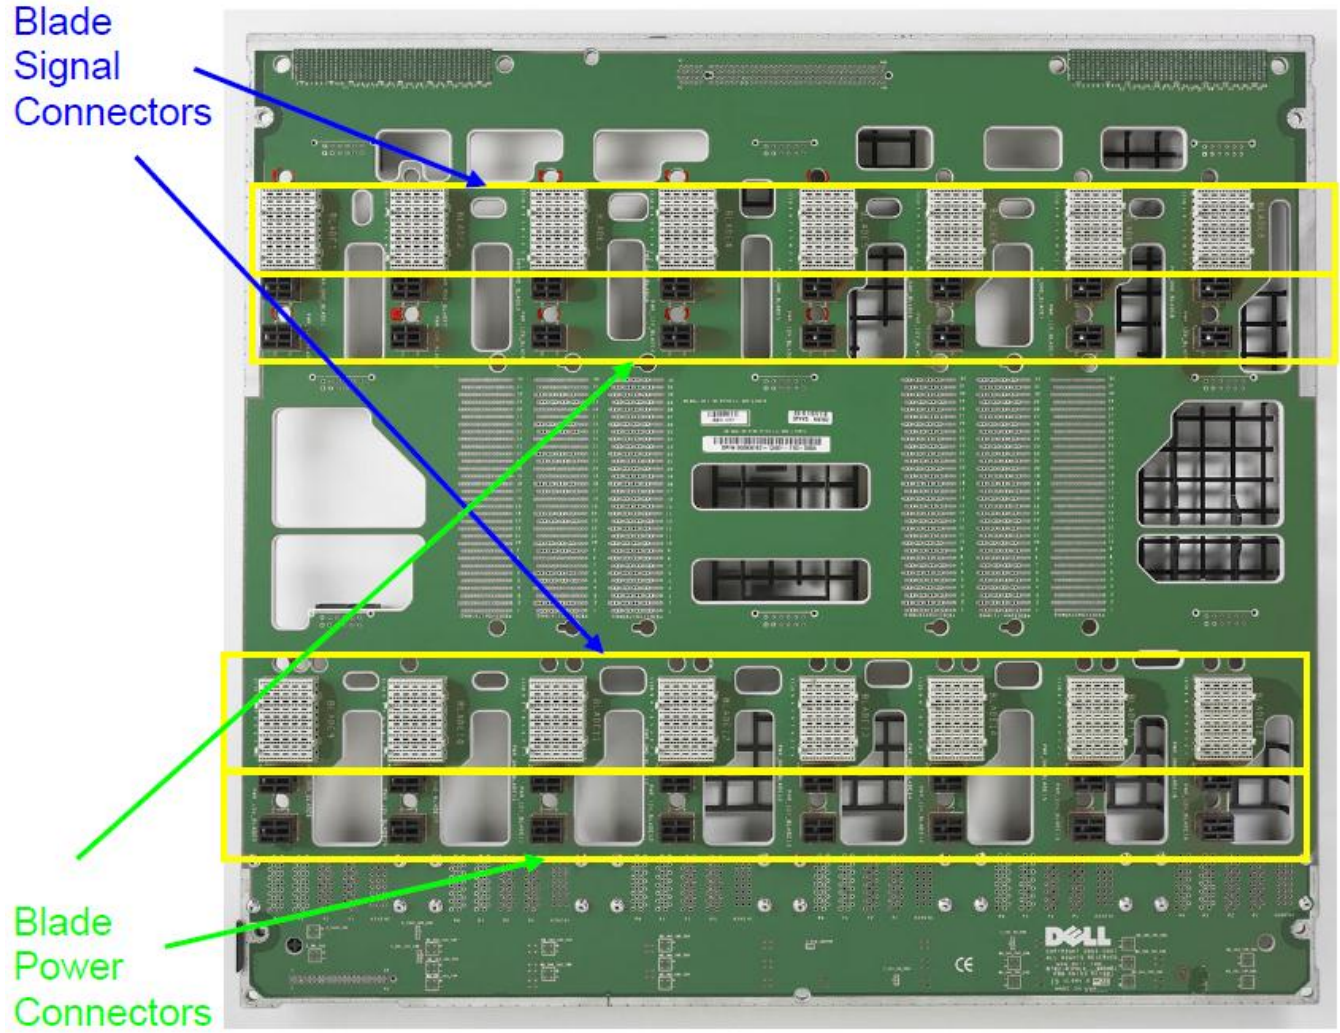

Figure 27. M1000e Midplane Front View

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. The midplane has been carefully designed to minimize the impact to the overall system airflow.

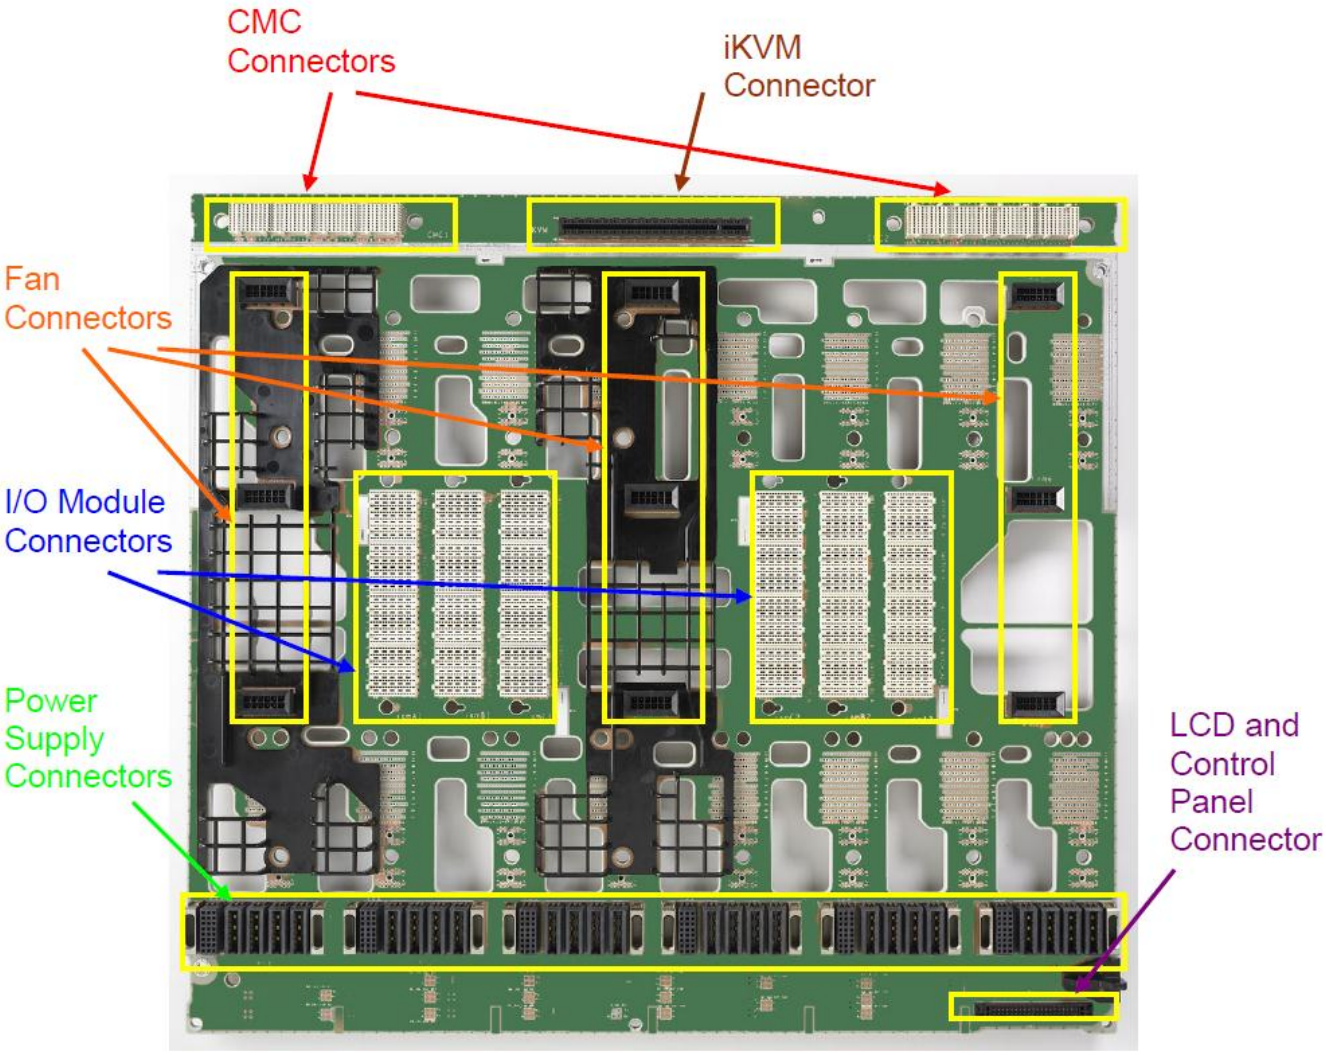

Figure 28. M1000e Midplane Back View

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

8 Embedded NICs/LAN on Motherboard (LOM)

For information on embedded NICs/LOMs, see the *Dell PowerEdge Mxxx Technical Guide* on Dell.com for each of the compatible blade servers.

9 Input/Output (I/O)

9.1 Overview

The Dell™ PowerEdge™ M-series provides 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. Customers may 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. See the I/O Connectivity section in the About Your System chapter of the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on Support.Dell.com/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 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
- 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 through the midplane to the associated I/O Modules (IOM) in the back of the enclosure, which then connect to the customer's 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 may 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, 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 29.

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.

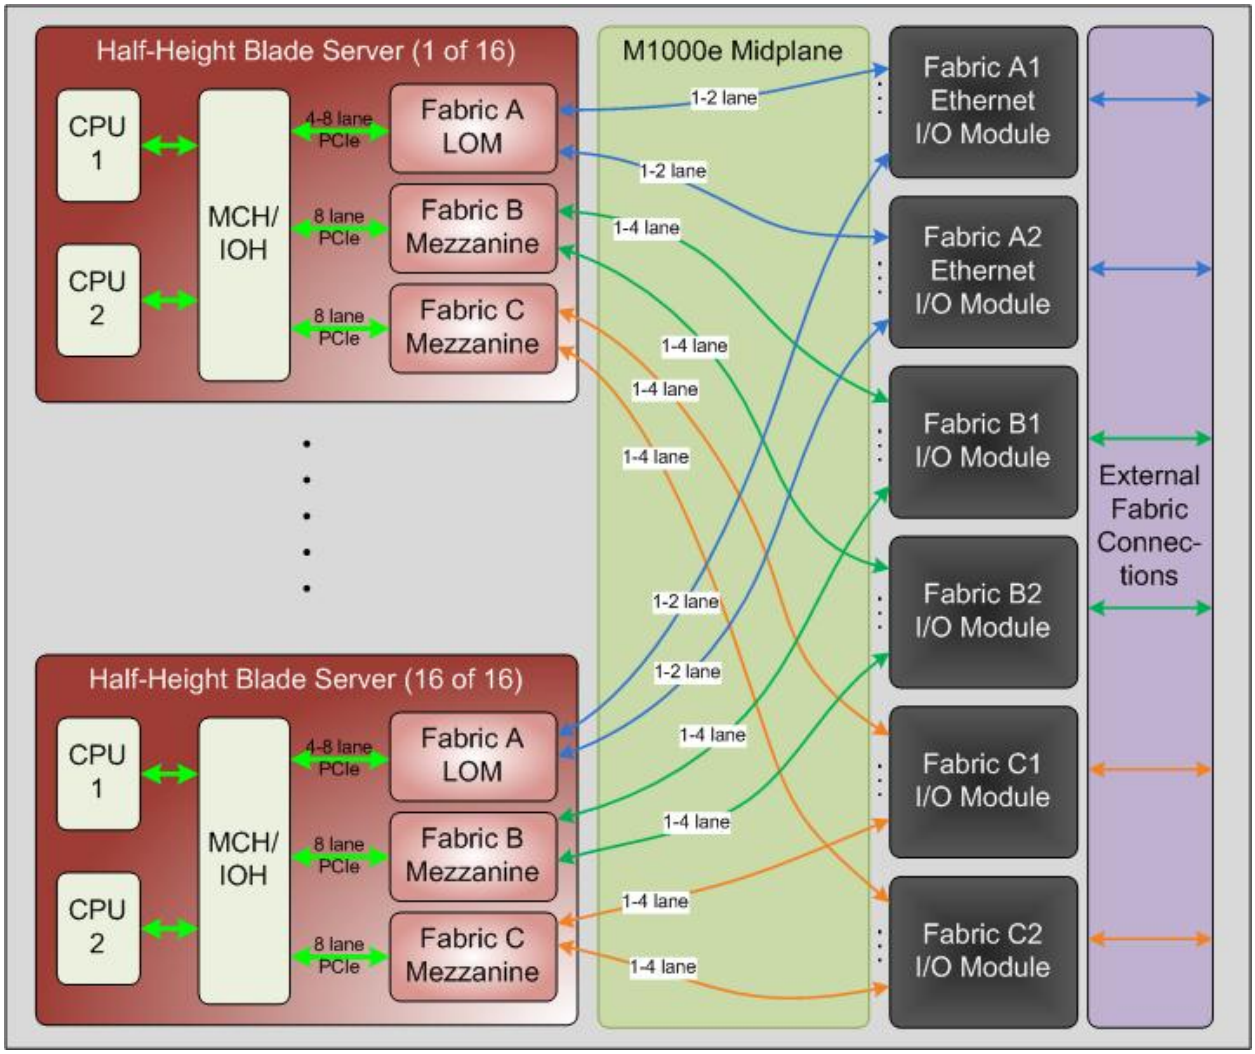

Figure 29. High Speed I/O Architecture

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.

Figure 30. 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

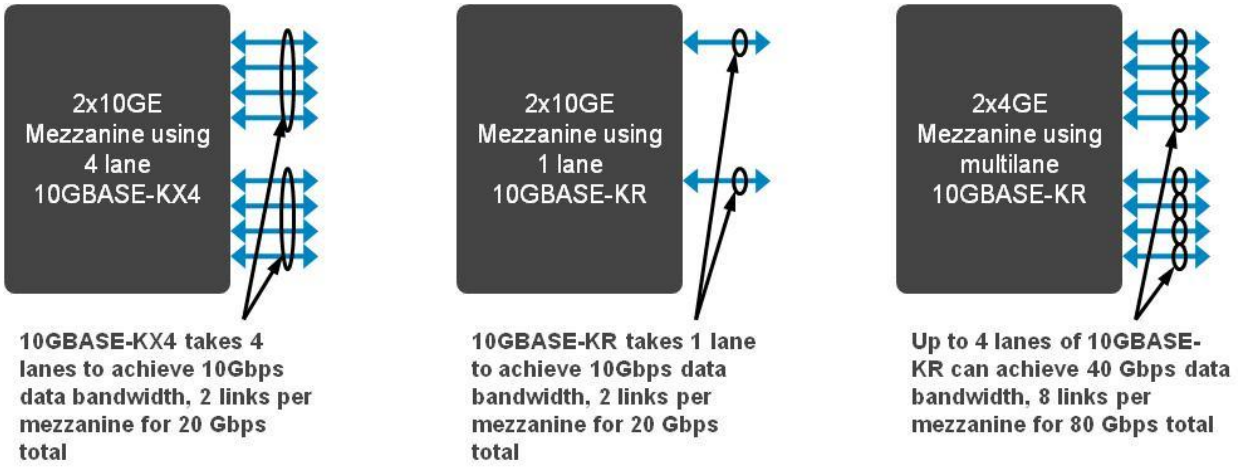

Figure 31. Ethernet Growth Path

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.

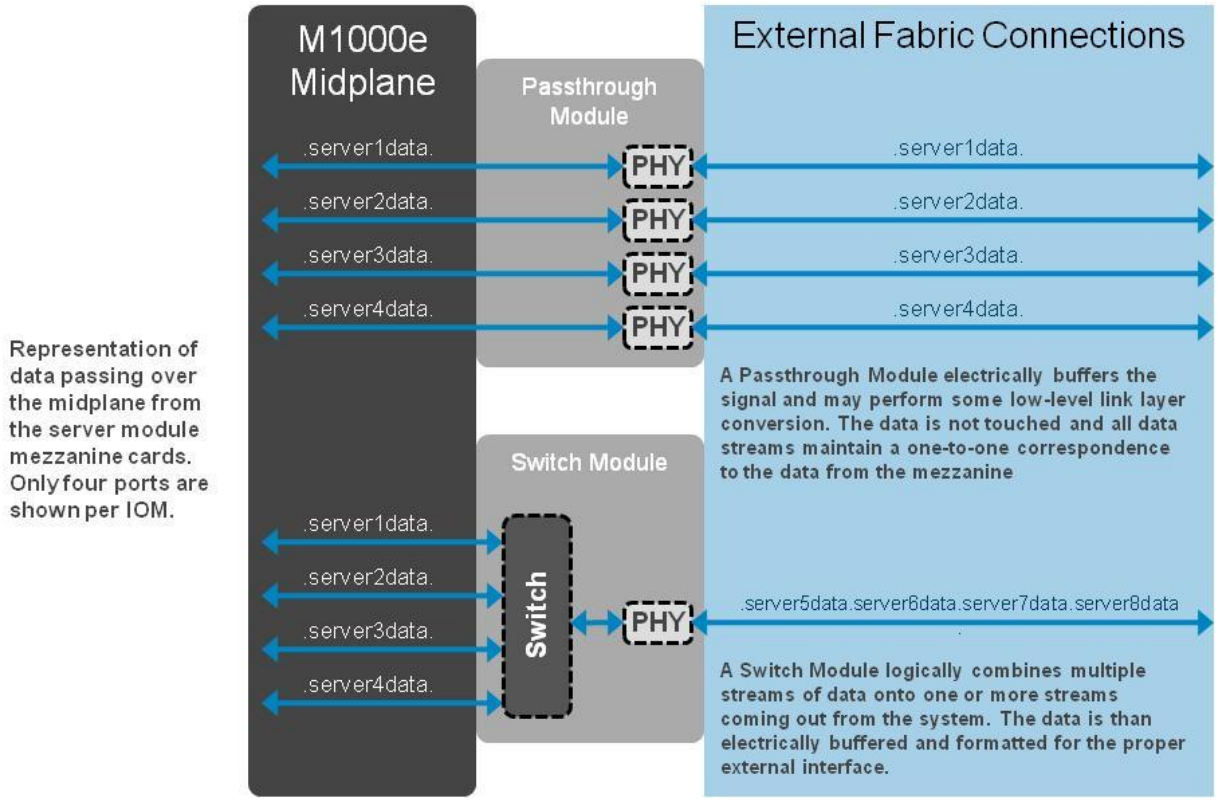

Figure 32. Difference Between Passthroughs and Switch Modules

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. Supported I/O modules include:

- PowerConnect M6220 switch; GbE + 10GbE uplinks and stacking
- PowerConnect M6348 switch; 48 1GbE ports + 10GbE uplinks
- PowerConnect M8024 10Gb Ethernet switch (SFP+, CX4, and 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 and stacking
- Cisco 3130g and 3130x switches can be combined in a stack
- 1Gb Ethernet pass-through module
- 10Gb Ethernet pass-through module (SFP+)
- Brocade® 8Gb fibre channel switch
- 4Gb Fibre channel pass-through

- Mellanox® DDR (20Gb) InfiniBand switch
- Mellanox QDR (40Gb) InfiniBand switch

For an up-to-date list of supported I/O modules, 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 Connectivity section in the About Your System chapter in the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on Support.Dell.com/Manuals.

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 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.

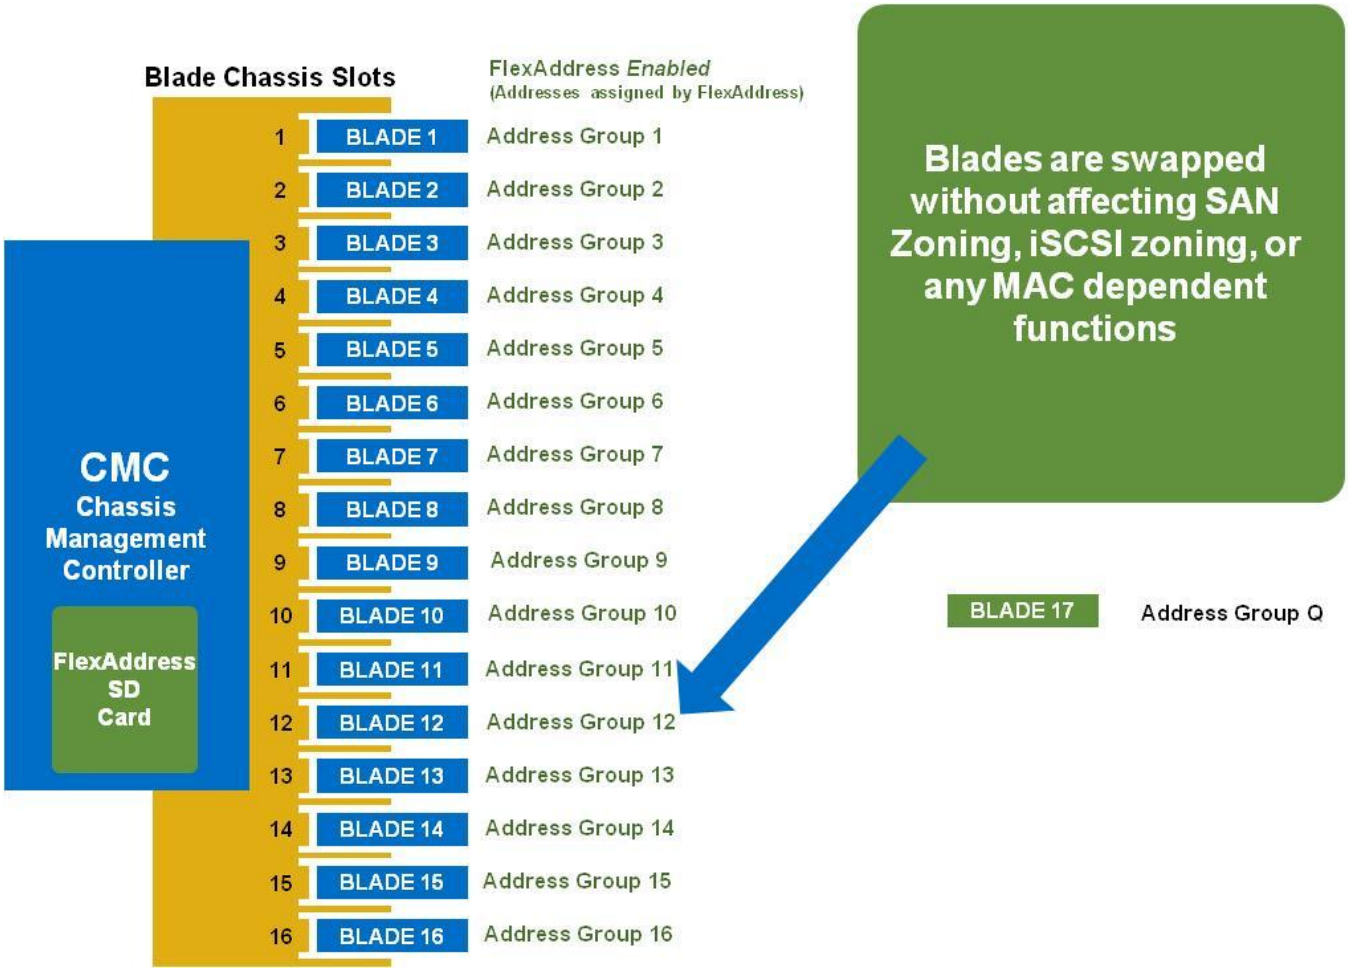

Figure 33. FlexAddress Addresses

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.

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, 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

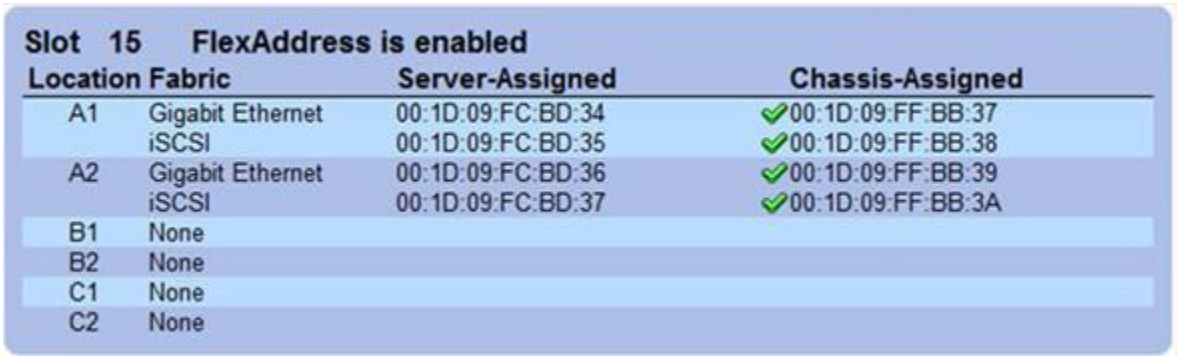

Figure 34. FlexAddress Screen in the CMC

Table 13. 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 and 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.

For information on implementing FlexAddress on an M1000e, see the Using FlexAddress chapter in the *Dell Chassis Management Controller Firmware x.x User Guide* on Support.Dell.com/Manuals.

10 Storage

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

Figure 35. Examples of Major Storage Platforms Supported

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

For details, see the *Dell PowerEdge Mxxx Technical Guide* for each of the compatible blade servers offered on Dell.com.

11 Video

The iKVM supports a video display resolution range from 640 x 480 at 60 Hz up to 1280 x 1024 x 65,000 colors (noninterlaced) at 75 Hz.

Table 14. Supported Video Modes

Resolution	Refresh Rate (Hz)	Color Depth (bit)
640 x 480	60, 72, 75, 85	8, 16, 32
800 x 600	56, 60, 72, 75, 85	8, 16, 32
1024 x 768	60, 72, 75, 85	8, 16, 32
1152 x 864	75	8, 16, 32
1280 x 1024	60, 75, 85	8, 16
1280 x 1024	60	32

12 Rack Information

12.1 Overview

The Dell™ RapidRails™ static rail system for the PowerEdge™ M1000e provides toolless support for racks with square mounting holes including all generations of Dell racks except for the 4200 and 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 toolless installation in 4-post racks with square mounting holes.

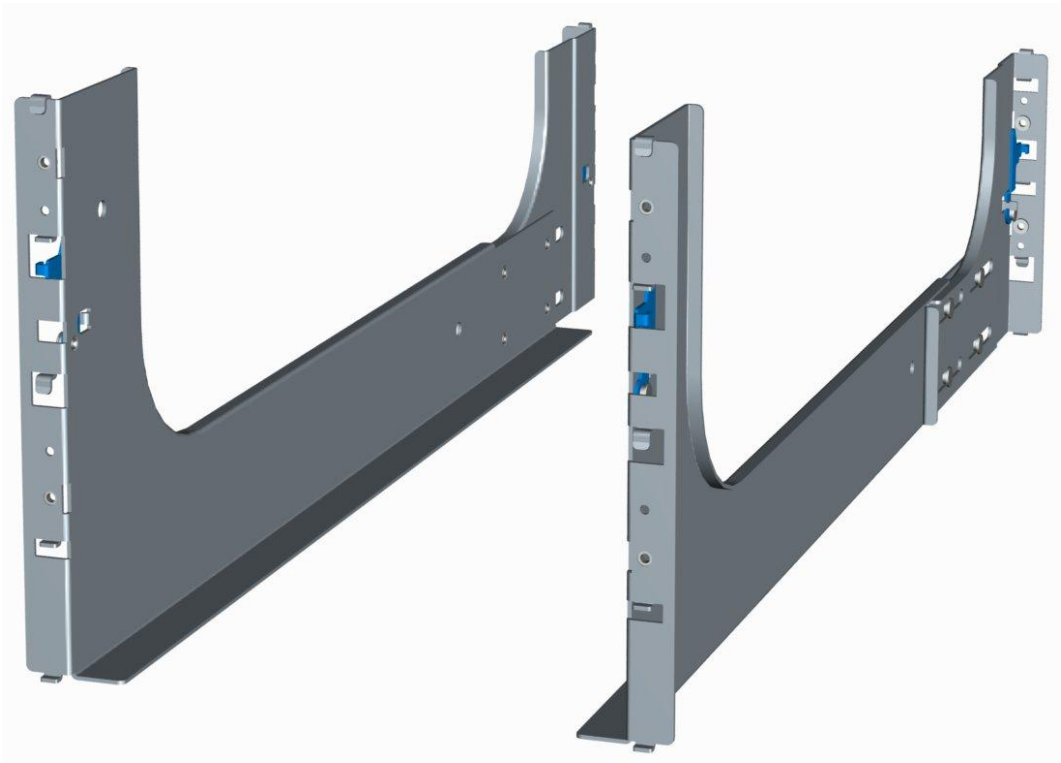

Figure 36. M1000e RapidRails Static Rails

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.

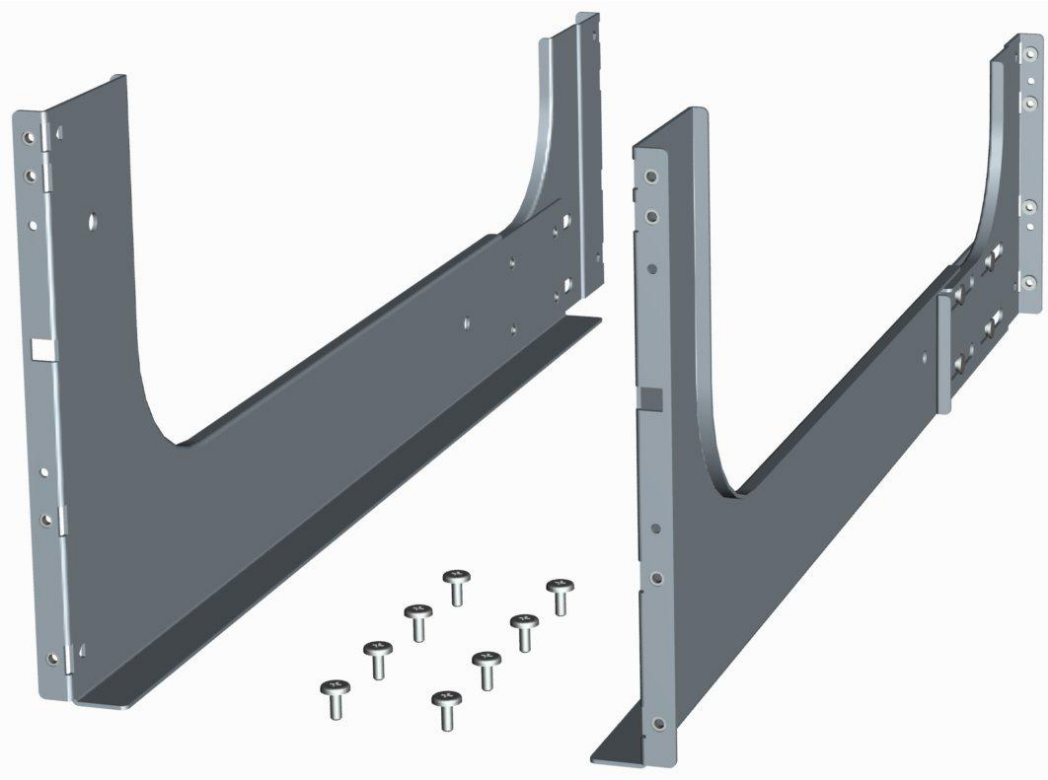

Figure 37. M1000e VersaRails Static Rails

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.

Table 15. Supported 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	√	X	X	X	X	712	755	—	—	—	—
VersaRails	Static	√	√	X	X	X	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 38. 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 39. 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.

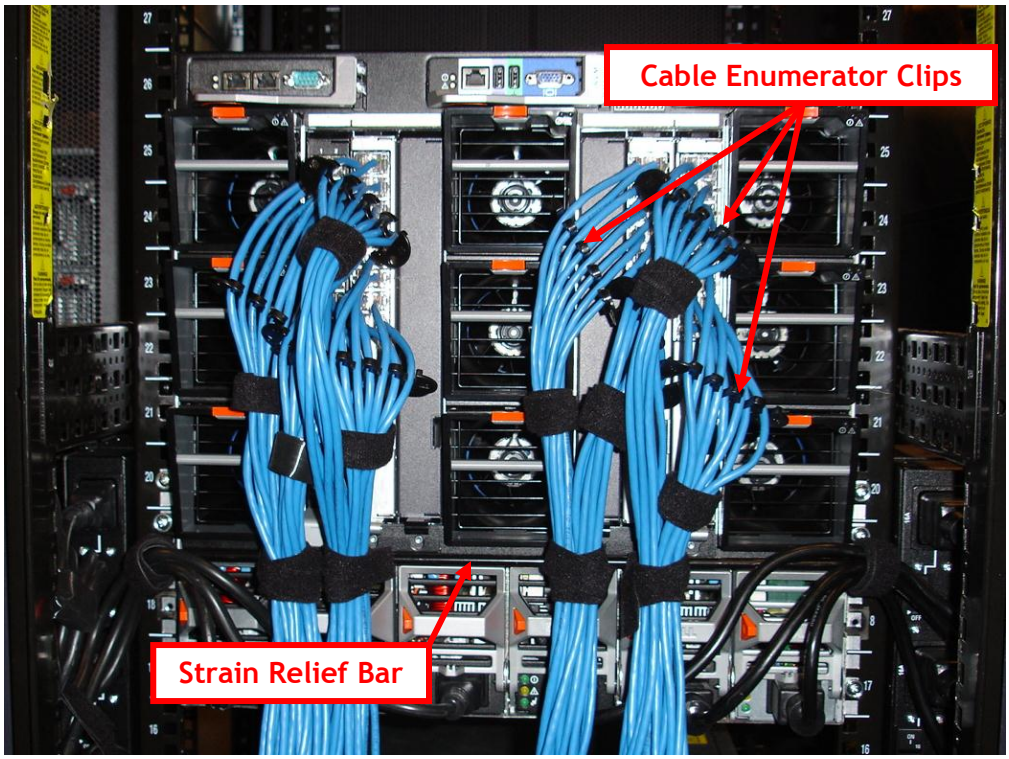

Figure 40. M1000e Strain Relief Bar and Cable Enumerator Clips

More information can be found in the *Rack Installation Guide* on Support.Dell.com/Manuals.

13 Virtualization

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

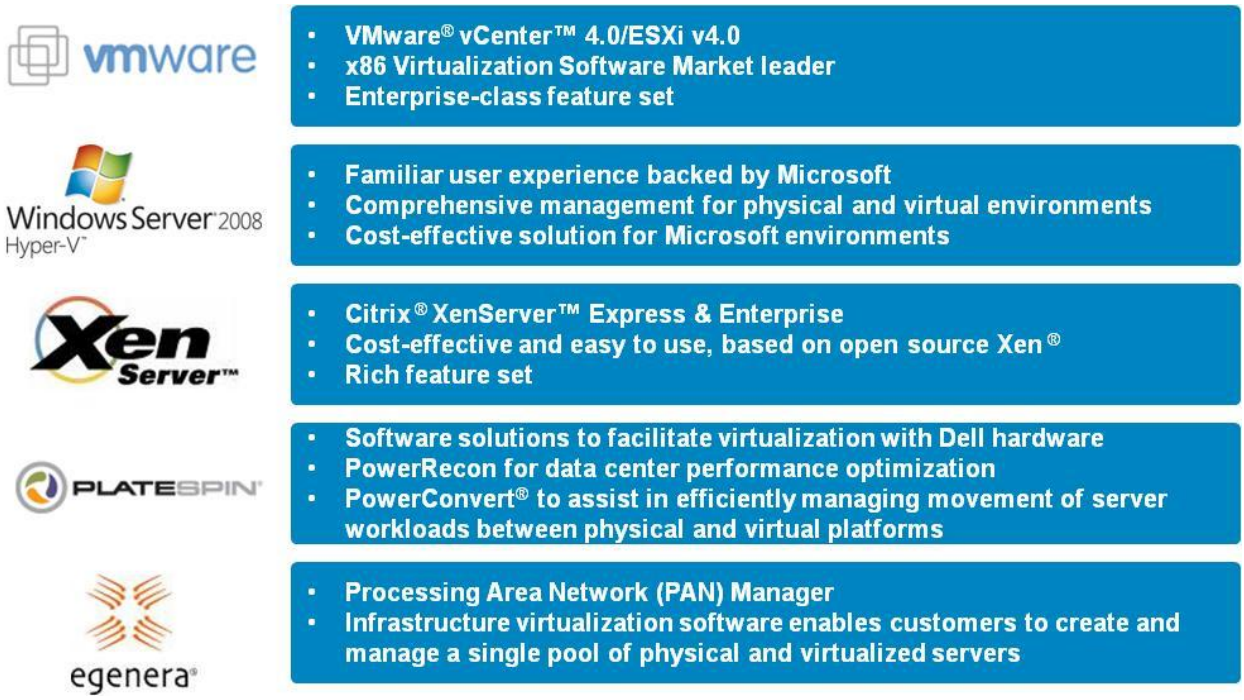

Figure 41. Examples of Major Virtualization Platforms Supported

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-swappable drives
- Power and cooling efficiency
  - Super-efficient power supplies
  - Optimized airflow
  - Best-in-class fan technology

14 Systems Management

14.1 Overview

The Dell™ PowerEdge™ M1000e server solution offers a holistic management solution designed to fit into any customer data center. The solution offers the following 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, 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 Integrated KVM (iKVM) module for local KVM access. Full USB access is available through the server module front panel.

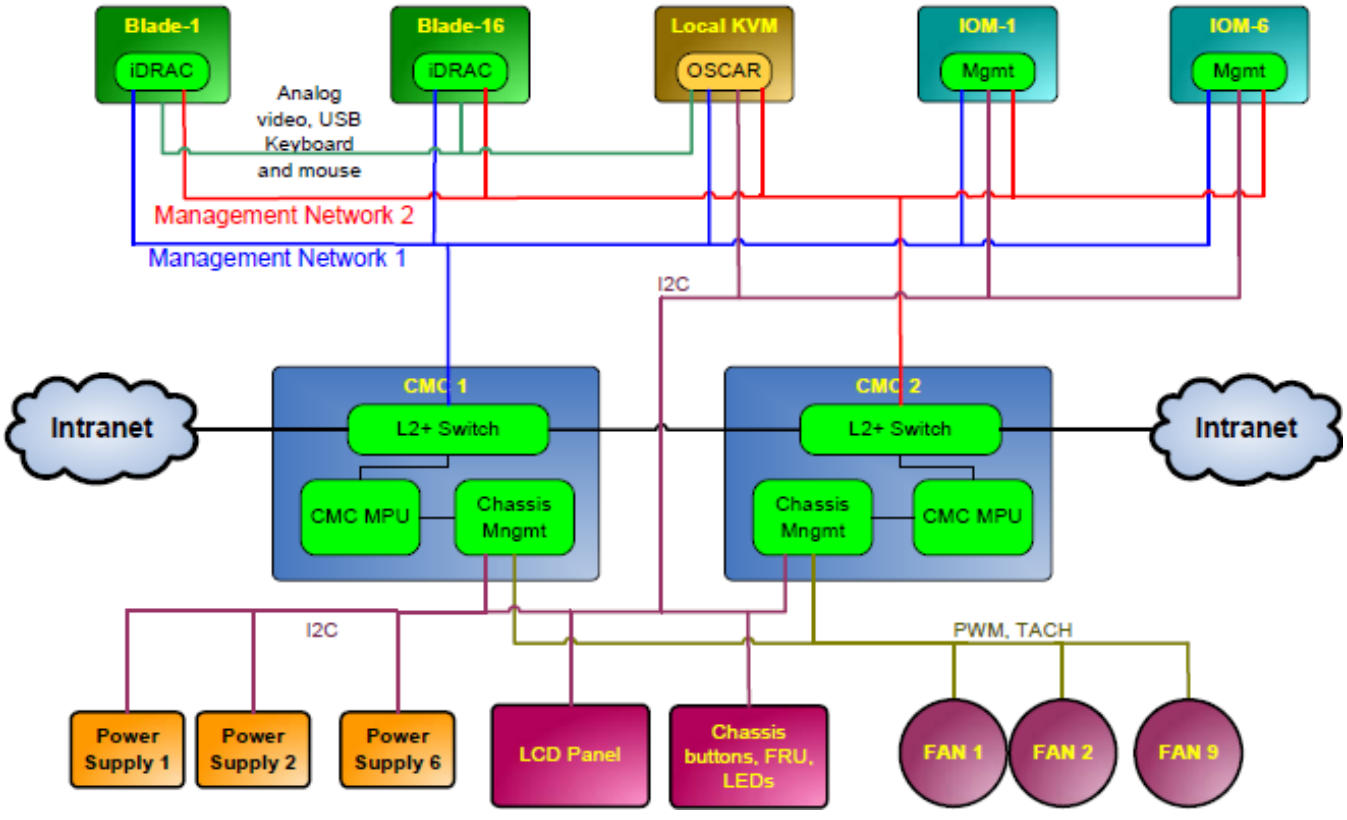

Figure 42. System Management Architecture Diagram

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
- Virtual Media
- Virtual KVM
- IPMI 2.0 Out Of Band management
- Serial over LAN redirection
- SMASH CLP
- Blade status and inventory
- Active power management
- Integration with Active Directory
- Security, Local and Active Directory

Traditional IPMI based Baseboard Management Controller (BMC) features like hardware monitoring and power control are supported.

The Lifecycle controller on eleventh-generation PowerEdge servers offers additional features including the following:

- Consolidated interface for OS install, hardware configuration, updates, and diagnostics
  - Reduces task time and speed of deployment
  - Eliminates physical 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 lifecycle log includes ability to add entries
- Out-of-Band configuration/update: With the introduction of Lifecycle Controller (LC) 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 or greater supports updates with BIOS, diagnostics, driver pack, USC, RAID controller firmware, iDRAC firmware, and NIC firmware.
- CMC 3.0 or greater will enable these updates by one-to-many using the CMC GUI/CLI.

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
  - 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

- 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
  - 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 integration
    - OpenLDAP coming in CMC 3.0
  - 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

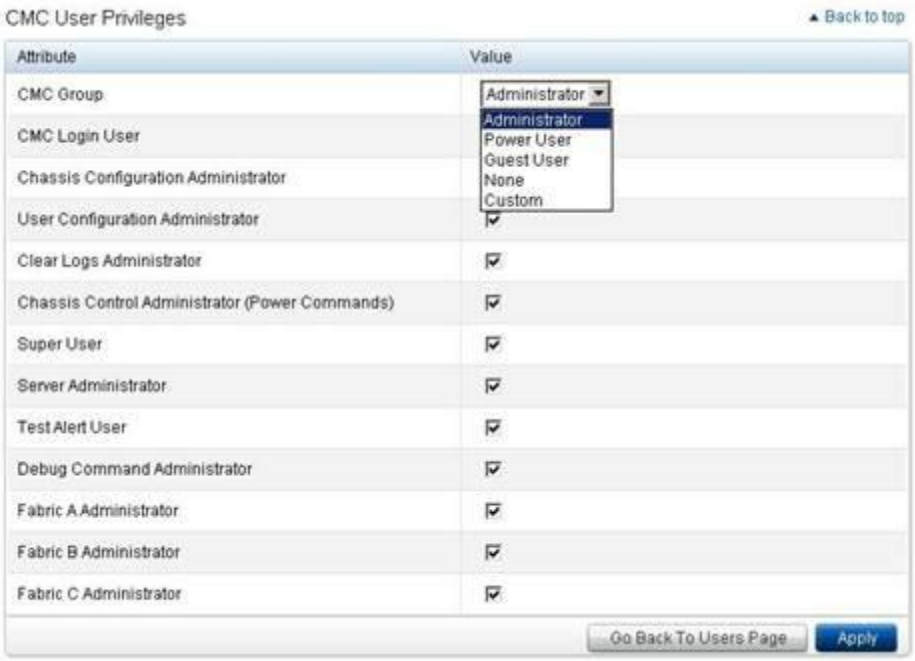

- Support for industry-standard interfaces
  - SMASH CLP
  - WSMAN
  - CIM XML
  - SNMP

The Integrated Dell Remote Access Controller (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 43. 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.

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 CMC Module section in the About Your System chapter in the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on on Support.Dell.com/Manuals.

14.4 Integrated Keyboard and Mouse Controller (iKVM)

Figure 44. M1000e iKVM

The M1000e modular enclosure supports one optional Integrated KVM (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 iKVM Module section in the About Your System chapter in the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on on Support.Dell.com/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 45. 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\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.

For more information on the iKVM see the *Dell PowerEdge Modular Systems Hardware Owner's Manual and CMC User Guide* on Support.Dell.com/Manuals, and at Dell.Avocent.com.

15 Peripherals

Common peripherals for the Dell™ PowerEdge™ M1000e include the following:

- An external USB DVD-ROM Drive, 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 (UPS) for racks, which provide 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
    - 30A for a medium to lightly loaded chassis
    - 60A for a heavily loaded chassis
  - For 3 phase:
    - 30A three phase for a heavily loaded single chassis
    - 50A or 60A 3 phase for multiple chassis
  - Generally, two 3-phase circuits are run to the rack through a PDU, then the PDU breaks out single phases to each PSU.

Appendix A. Regulatory Certifications

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

Appendix B. Status Messages

A 1. LCD Status Messages

For a detailed list of LCD status messages, see the LCD Status Messages section in the About Your System chapter in the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on on Support.Dell.com/Manuals.

A 2. System Status Messages

For a detailed list of system status message, see the System Messages section in the About Your System chapter in the *Dell PowerEdge Modular Systems Hardware Owner's Manual* on on Support.Dell.com/Manuals.

Appendix C. Additional Information

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

The *Dell PowerEdge M1000e Configuration Guide*, *Dell PowerEdge Modular Systems Hardware Owner's Manual*, and *CMC Administrator Reference Guide* or *CMC User's Guide* on Support.Dell.com/Manuals each contain a wealth of additional information about the PowerEdge M1000e's capabilities.