

Scaling Direct Attach Storage

with the Dell PowerVault MD1000 Storage Expansion Enclosure

Serial Attached SCSI (SAS) introduces new levels of performance and scalability for direct attach storage. To take advantage of industry-standard SAS technology, Dell offers the Dell™ PowerVault™ MD1000 storage expansion enclosure. When combined with the Dell PowerEdge™ RAID Controller 5/Extended and the Dell OpenManage™ Server Administrator application, the PowerVault MD1000 can provide a comprehensive SAS storage system.

BY SANGRAM KADAM

Related Categories:

Dell PowerVault storage

RAID controllers

Serial ATA (SATA)

Serial Attached SCSI (SAS)

Storage

Visit www.dell.com/powersolutions
for the complete category index.

For more than 20 years, SCSI technology has been deployed to meet mainstream enterprise storage requirements. With the arrival of Serial Attached SCSI (SAS)—the next generation of SCSI-based disk technology—enterprises can deploy storage systems that are more flexible and scalable than ever before.

As the industry's first 3.5-inch SAS-based disk enclosure, the PowerVault MD1000 storage expansion enclosure builds upon the Dell tradition of offering expandable server storage. By incorporating SAS technology, the PowerVault MD1000 is designed to offer a point-to-point architecture capable of delivering up to 12 Gbps performance. Along with the PCI Express-based Dell PowerEdge RAID Controller 5/Extended (PERC 5/E), the PowerVault MD1000 can help enterprises to meet their growing direct attach storage needs.

Development of SCSI technology

Over the years, parallel SCSI has matured and become the storage interconnect of choice for many mainstream

enterprises deploying server-centric storage. Although parallel SCSI has continued to mature with Ultra320 SCSI—the latest version—the SCSI standard has reached performance saturation because of the limitations of the parallel bus architecture. Also, parallel SCSI can limit scalability because of the restricted addressability of SCSI devices, which is 16 IDs per SCSI bus.

The SAS protocol was developed to address needs for high performance and addressability of direct attach storage. This protocol builds upon mature parallel SCSI technology and implements point-to-point architecture that enables current speeds of 3 Gbps per connection—thus, a 4X connection provides up to 12 Gbps. The SAS protocol also provides support for SAS and Serial ATA (SATA) devices in the same environment.

With its point-to-point architecture, the SAS protocol implements expanders to allow the connection of more drives than ports. SAS expanders are physical device-level,

store-and-forward switches that let a single SAS port address more than one device.

A SAS domain is formed by a network of SAS devices that are interconnected by SAS expanders. Any device in a SAS domain can

	SCSI	SAS
Current performance capability	320 MB/sec	3 Gbps*
Projected performance capability	320 MB/sec	6 Gbps*
Cable length	12 meters	6 meters
Scalability	16 devices per bus	16,000-plus devices per domain
Architecture	Parallel bus	Point to point
Device compatibility	SCSI devices	SAS and SATA devices

*SAS performance is for a 1X point-to-point connection.

Figure 1. Comparison of SCSI and SAS technologies

Feature	PowerVault 220S/221S	PowerVault MD1000
Rack height	3U	3U
Number of HDD bays	Fourteen 3.5-inch drives	Fifteen 3.5-inch drives
Manageability	In-band SCSI	In-band SCSI
Backplane options	Joined or split bus	Unified or split bus
Drive support	Ultra160, Ultra320, Low Voltage Differential SCSI	SAS 3 Gbps; SATA II
SCSI and RAID cards	PERC 2 PERC 3/Dual Channel (DC) PERC 3/Dual Channel Light (DCL) PERC 3/Quad Channel (QC) PERC 4/Single Channel (SC) PERC 4/DC PERC 4/DC integrated (Di) Adaptec 39160 host bus adapter	PERC 5/E
Cluster support	Yes	No
Expansion	No	Yes, daisy-chain up to three enclosures per connection
Storage management software	Dell OpenManage Array Manager and Dell OpenManage Server Administrator	Dell OpenManage Server Administrator
Rack or stand-alone tower	Both	Rack only
Drive hot-plug support	Yes	Yes
Hot-plug fans/blowers	Yes	Yes, via the power supply-cooling fan module
Hot-plug power supplies	Yes	Yes, via the power supply-cooling fan module
Enclosure management configuration	Redundant or nonredundant	Redundant or nonredundant
Power supply configuration	Redundant or nonredundant	Redundant
Fans/blowers	Redundant	Redundant

Figure 2. Comparison of PowerVault 220S/221S and PowerVault MD1000 disk enclosures

communicate with any other device in the same SAS domain. And unlike SCSI, SAS enables bandwidth performance to scale as more devices are added to the domain. Figure 1 provides a comparison of SCSI and SAS capabilities.

Dell SAS implementation

The Dell SAS implementation comprises both hardware and software components. The hardware components consist of the Dell PowerVault MD1000 enclosure and the PCI Express-based PERC 5/E RAID controller. The software component consists of Dell OpenManage Server Administrator software.

Dell PowerVault MD1000 external disk enclosure

The Dell PowerVault MD1000 is the external SAS disk enclosure. Redundancy is built into the enclosure with optional dual enclosure management modules (EMMs) and dual power supply-cooling fan modules. With dual EMMs, control and monitoring of the enclosure elements continue even if one EMM fails. This allows the functional EMM to manage the enclosure—monitoring and controlling the alarms, LEDs, power supplies, and fans.

Each power supply-cooling fan module in the PowerVault MD1000 contains an integrated power supply with two fans. Although the enclosure can sustain the loss of one power supply, three of the four fans must be functional to maintain proper cooling. During a hot replacement, cooling is supported for a maximum of five minutes when only one power supply-cooling fan module is present in the enclosure. The remaining fans will speed up to provide the necessary cooling.

The PowerVault MD1000 provides storage provisioning flexibility. This is achieved by two backplane modes—unified and split—offered in the enclosure. Administrators can easily set the backplane mode by using the toggle switch located on the front of the PowerVault MD1000 and power cycling the enclosure for the change to take effect.

Unified mode lets administrators access all 15 hard disk drives (HDDs) from a single controller and provides the ability to daisy-chain up to two additional enclosures. This allows access to a total of 45 HDDs with a maximum of 13.5 TB (45 × 300 GB, given the currently supported drive sizes). Split mode lets two controllers share the same enclosure in an eight-to-seven HDD split—that is, one controller has access to eight of the HDDs, while another controller has access to the other seven HDDs. An enclosure in split configuration does not permit daisy-chaining additional enclosures.

These PowerVault MD1000 features are designed to provide cost-effective, scalable storage with SAS or SATA drives. Figure 2 provides a features comparison of the PowerVault MD1000 storage expansion enclosure and the previous-generation PowerVault 220S/221S SCSI disk storage enclosure.

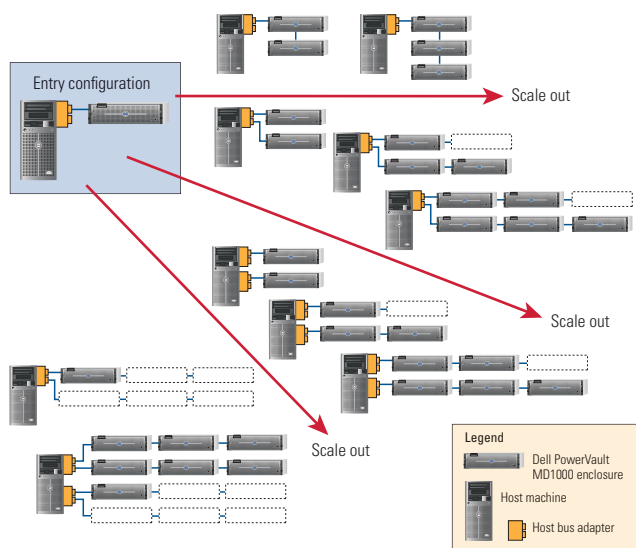


Figure 3. High-capacity Dell PowerVault MD1000 SAS configuration

Dell PowerEdge RAID Controller 5/E

The PCI Express–based PERC 5/E provides RAID capabilities for the PowerVault MD1000 SAS enclosure. The RAID levels offered by this controller include RAID-0, RAID-1, RAID-5, RAID-10, and RAID-50. The controller has a 256 MB transferable cache and battery backup unit to help maintain data consistency in the event of a power loss. Besides basic RAID functions, the controller also supports the following features:

- Drive migration and roaming
- SAS device and topology discovery
- Persistent device mapping
- RAID capacity expansion
- RAID-level migration
- Background Patrol Read
- Self-Monitoring, Analysis, and Reporting Technology (SMART) error reporting

The PERC 5/E provides two 4X ports to connect the controller to the enclosure, providing 12 Gbps per port. The PowerVault MD1000 connects to the PERC 5/E 4X channel with industry-standard 4X SAS cables.

Dell management software

Offline/preboot management and configuration of the PowerVault MD1000 can be accomplished with the BIOS Configuration Utility, which is launched by pressing Ctrl + R immediately after the power-on self-test (POST) of the attached server. Remote preboot configuration and management of the storage subsystem are facilitated by support for console redirection. Navigation keys are

designed to follow the key functionality of Microsoft® Windows® Explorer—the arrow and Tab keys are the main keys for in-screen movement, and Ctrl + N (next) and Ctrl + P (previous) allow changing between the Virtual Disk Management, Physical Disk Management, and Controller Management screens.

Online monitoring and managing of the PowerVault MD1000 can be easily accomplished by using Dell OpenManage Server Administrator. This application provides administrators with a graphical user interface to access the PowerVault MD1000 over a supported Web browser (Microsoft Internet Explorer for Microsoft Windows platforms or Mozilla Firefox for Linux® platforms) from a remote or local station.

Deployment options for the PowerVault MD1000 SAS implementation

The Dell PowerVault MD1000 disk enclosure can be tailored to meet different business needs. This section provides two examples of how this enclosure can be used to respond to capacity and efficiency requirements.

Gaining high storage capacity

To obtain optimal storage capacity using a PowerVault MD1000–based configuration for a single server, the unified mode of the enclosure should be selected. This allows the server to access all 45 drives in a three-enclosure cascading configuration from a single PERC 5/E connector. Three additional enclosures can be connected to the other connector on the PERC 5/E controller, giving the server access to 90 HDDs in a fully populated configuration. This configuration lets administrators start with a single enclosure and scale out as storage needs grow, thus allowing for a consistent return on investment (see Figure 3).

Connecting to multiple servers

To efficiently leverage the storage enclosure, administrators should deploy the enclosure in split mode. This will permit two servers to access a portion of the enclosure. The first server connects the PERC 5/E to the primary EMM on the PowerVault MD1000, accessing seven drives. The second server connects the PERC 5/E to the secondary EMM on the same enclosure, accessing eight drives (see Figure 4). The allocated storage space is not shared.




Figure 4. Multi-server Dell PowerVault MD1000 SAS configuration

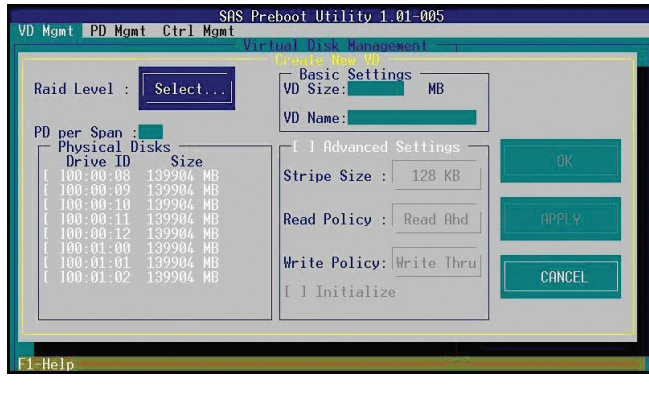


Figure 5. BIOS Configuration Utility

Management tools for the PowerVault MD1000 SAS implementation

After the hardware is physically deployed, virtual disks can be configured on the drives in the enclosure using either the BIOS Configuration Utility or Dell OpenManage Server Administrator.

BIOS Configuration Utility

The offline/preboot BIOS Configuration Utility is limited to managing the drives in the storage enclosure. This utility can help administrators manage virtual and physical disks as well as the RAID controller (see Figure 5).

The BIOS Configuration Utility provides virtual disk management, which includes the following capabilities:

- Viewing the controllers in the system and their properties
- Viewing a tree or list of all virtual disks, physical disks, and space allocation in the controllers
- Creating new virtual disks by specifying virtual disk names, RAID levels, sizes, physical drives, number of physical drives to span, and other parameters
- Configuring advanced features such as stripe size and read/write policy
- Initializing a virtual disk
- Resetting configurations

Physical disk management performed through this utility encompasses the following capabilities:

- Viewing all the hard drives in the storage subsystem
- Starting or stopping a drive rebuild
- Visually identifying drives

- Unaffiliating a drive directly or online
- Configuring a global hot spare

The BIOS Configuration Utility also provides controller management, which includes the following capabilities:

- Revising the controller BIOS and firmware
- Revising the BIOS Configuration Utility
- Enabling and disabling the controller BIOS, alarms, clustering, and BIOS halt-on-error during POST
- Selecting a virtual disk from which to boot
- Selecting controller defaults and resetting the controller configuration

Dell OpenManage Server Administrator

Dell OpenManage Server Administrator (see Figure 6) has a larger set of management features than the offline/preboot BIOS Configuration Utility, including management of the enclosure and battery options. Management features provided by Dell OpenManage Server Administrator include the following:

- **Controller configuration:** Managing rebuilds, performing background initializations and consistency checks, configuring reconstruction rates, resetting configurations, enabling and disabling alarms, exporting logs, and configuring and starting the Background Patrol Read tool
- **Configuration and management of virtual disks:** Performing consistency checks, assigning and unassigning

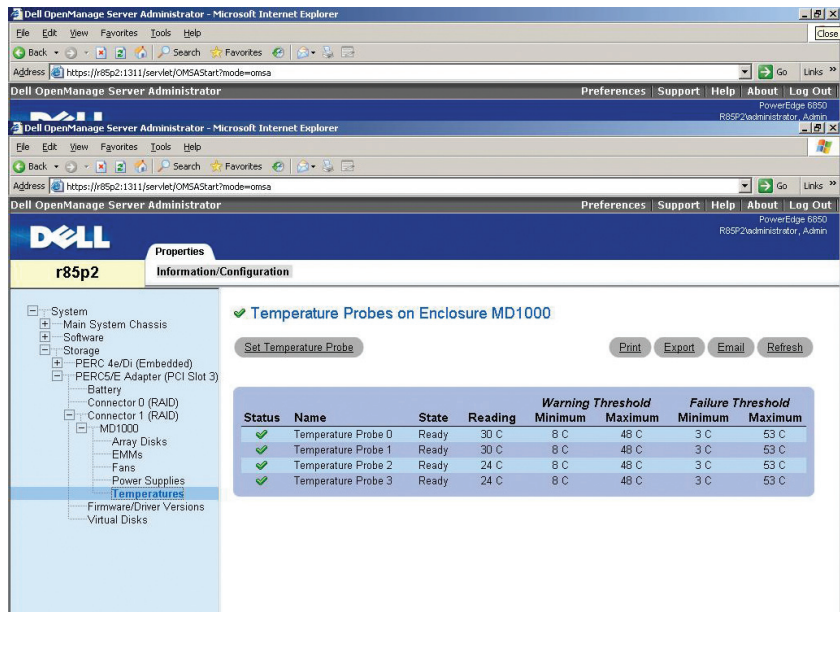


Figure 6. Enclosure information and configuration screen in Dell OpenManage Server Administrator

dedicated hot spares, identifying drives in virtual disks, changing read/write policies, starting slow or fast initializations, and reconfiguring virtual disks for capacity expansion or RAID-level migration

- **Battery configuration:** Starting and delaying learning cycles
- **Enclosure configuration:** Enabling and disabling alarms, setting asset information, identifying enclosures, configuring temperature thresholds, and performing firmware upgrades

The reconfiguration, capacity expansion, and RAID-level migration capabilities enable administrators to manage the expansion of their storage subsystems cost-effectively as required by the enterprise.

Furthermore, Dell OpenManage Server Administrator's management features facilitate monitoring the enclosure health and the status of the storage array subsystem. Administrators can perform hierarchical monitoring of the following conditions:

- **Controller health:** Battery and connector status and firmware revisions
- **Enclosure health:** Status of array disks, management modules, fans, power supplies, and temperature (see Figure 7)

Suboptimal conditions in any enclosure or storage array components trigger a notification that rises to the top of the monitoring hierarchy. For administrators who use Dell OpenManage IT Assistant (ITA) 7 to monitor the IT infrastructure, the notification displays on the ITA monitoring console. From ITA, administrators can then probe the lower tiers of the monitoring hierarchy—such as Dell OpenManage Server Administrator—to identify and resolve the issue from a remote management station.

SAS: The next step in server-centric storage

SCSI has been the primary technology for disk-based storage for more than 20 years. It has matured into a highly robust, reliable, effective protocol for enterprise storage subsystems. But it has also reached its limits of speed and scalability.

The Serial Attached SCSI standard for disk-based storage devices embraces the advantages of traditional SCSI while allowing for higher speeds and greater scalability than SCSI can provide. Together, the Dell PowerVault MD1000 storage expansion enclosure, the PERC 5/E controller, and Dell OpenManage Server Administrator can provide a SAS implementation that enables new levels of perfor-





Figure 7. Enclosure health screen in Dell OpenManage Server Administrator

mance and scalability for enterprises seeking to deploy server-centric storage. With the launch of these SAS-based products, Dell continues to support its pledge to help build scalable enterprises by offering data center components designed to simplify operations, improve utilization, and promote cost-effective scaling. 

Sangram Kadam is an engineer in the Dell Enterprise Quality–Storage Group. He has a B.S. in Computer Engineering from Bombay University and an M.S. in Computer Science from North Carolina State University.

FOR MORE INFORMATION

Dell PowerVault storage:

www.dell.com/powervault

SCSI Trade Association:

www.scsita.org

Dell OpenManage Server Administrator User's Guide:

support.dell.com/support/edocs/software/svradmin