

Exploring the Dell PowerVault MD3000 Storage System

BY MANJUSHA GOPAKUMAR
VISHAL KADAM
PAUL WINSTON

The Dell™ PowerVault™ MD3000 external RAID enclosure allows multiple servers to access a shared pool of expandable storage, helping enterprises avoid the costs and complexities of internal storage. This article outlines the hardware, configuration, and management features of this storage system.



The Dell PowerVault MD3000 storage array is designed to provide flexible, high-performance, highly scalable RAID storage without compromising data integrity. Its integration of two high-performance RAID controllers, Dell Modular Disk Storage Manager software, and automatic event monitoring helps reduce total cost of ownership and maximize data protection.

The PowerVault MD3000 external RAID array can support up to forty-five 3.0 Gbps Serial Attached SCSI (SAS) disks and includes redundant power supplies and fans. Connectivity between host servers and the storage is provided by two Dell SAS 5/E host bus adapters (HBAs). The PowerVault MD3000 is designed primarily for high-performance two-node clusters; however, its dual-port RAID controllers can also provide shared nonredundant storage for up to four servers. Other key features include the following:

- Mirrored data cache of up to 512 MB on each RAID controller
- In-band and out-of-band management using Dell Modular Disk Storage Manager
- Redundant, hot-swappable components and virtual disk failover
- Transparent failover for clustered operating systems to help maximize data protection
- Virtual disk snapshot and virtual disk copy creation
- RAID level migration, capacity expansion, consistency

checks, and Self-Monitoring, Analysis, and Reporting Technology (SMART)

Figure 1 compares PowerVault MD3000 features with those of the PowerVault MD1000.

Configuration options for the Dell PowerVault MD3000

Administrators can configure the PowerVault MD3000 as part of a nonredundant configuration or a redundant cluster configuration. Nonredundant configurations are recommended only for noncritical data storage; this type of configuration provides only a single data path, meaning a failed or removed cable, HBA, or RAID controller could cause the storage path to fail.

Redundant cluster configurations, in contrast, provide high availability for critical data and depend primarily on physical connectivity and the host multipath software. If a component fails or an error occurs on the data path to the primary RAID controller, the multipath driver—Microsoft® Multipath I/O (MPIO) on Microsoft Windows® operating systems and Multipath Proxy (MPP) on Linux® operating systems—fails over virtual disk ownership to the secondary RAID controller for I/O processing, helping ensure host data access.

The PowerVault MD3000 supports two redundant cluster configurations:

- **Single port:** Administrators can use this mode with a single storage path and single-port RAID controllers (see

Related Categories:

Backup, recovery, and archiving (BURA)

Dell PowerVault storage

Direct attach storage (DAS)

Serial Attached SCSI (SAS)

Storage

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

	PowerVault MD1000	PowerVault MD3000
Drive technology support	SAS and SATA	SAS
SMART support	SAS and SATA	SAS
RAID support	RAID-0, RAID-1, RAID-5, RAID-10, and RAID-50	RAID-0, RAID-1, RAID-5, and RAID-10
RAID level migration	Limited	Any supported RAID level to any other supported RAID level
Host support	Maximum of 2	Maximum of 4
Management software	Dell OpenManage	Dell Modular Disk Storage Manager
Background operation priority	One per operation	One per virtual disk
Hot spare association	Global and dedicated	Global
Cache mirroring	N/A	✓
Disk group defragmentation		✓
Redundant controller support		✓
Redundant access to disk drives		✓
Virtual disk snapshots		✓
Virtual disk copies		✓
Cluster support		Microsoft Cluster Service and Oracle® Real Application Clusters

Figure 1. Comparison of Dell PowerVault MD1000 and PowerVault MD3000 storage

Figure 2). The Automatic Virtual Disk Transfer feature of Dell Modular Disk Storage Manager provides failover and is designed specifically for single-port cluster failover; it is automatically selected by host type during host access configuration.

- **Dual port:** Administrators can use this mode with dual storage paths and with dual-port RAID controllers (see Figure 3, which shows a configuration utilizing two HBAs on each

host server). Multipath drivers such as MPIO and MPP are installed on host servers and provide I/O path failover. The controller firmware executes ownership transfer operations for virtual disks and owns the logical units (LUNs).

Management and monitoring features

Administrators can manage and monitor the PowerVault MD3000 using the Dell Modular

Disk Storage Manager graphical user interface (GUI) or command-line interface (CLI). The GUI includes wizard-based tools and a task-based structure designed to simplify installation, configuration, management, and diagnostic tasks. Modular Disk Storage Manager enables administrators to discover storage systems through either Serial Attached SCSI (SAS) or Ethernet, troubleshoot using diagnostic utilities, and manage virtual disks, hot-spare drives, firmware upgrades, and event logs (see Figure 4). It includes both in-band and out-of-band management capabilities, along with an optional event monitoring service (used to send alerts when a critical problem occurs).

In-band management is available when the PowerVault MD3000 is connected to the host server using SAS and Modular Disk Storage Manager is fully installed, including the host-based storage agent, multipath drivers, and management station package. Out-of-band management is available when client systems can connect to the PowerVault MD3000 using Ethernet and the Modular Disk Storage Manager management station package is installed.

The Modular Disk Storage Manager CLI gives administrators direct access to a script engine utility that reads commands and runs script files. Tasks administrators can perform using the CLI include the following:

- Create script command batch files when installing the same configuration on multiple storage arrays
- Run script commands on a storage array directly connected to a host, a storage array




Figure 2. Single-port cluster configuration using Dell PowerVault MD3000 storage




Figure 3. Dual-port cluster configuration using Dell PowerVault MD3000 storage

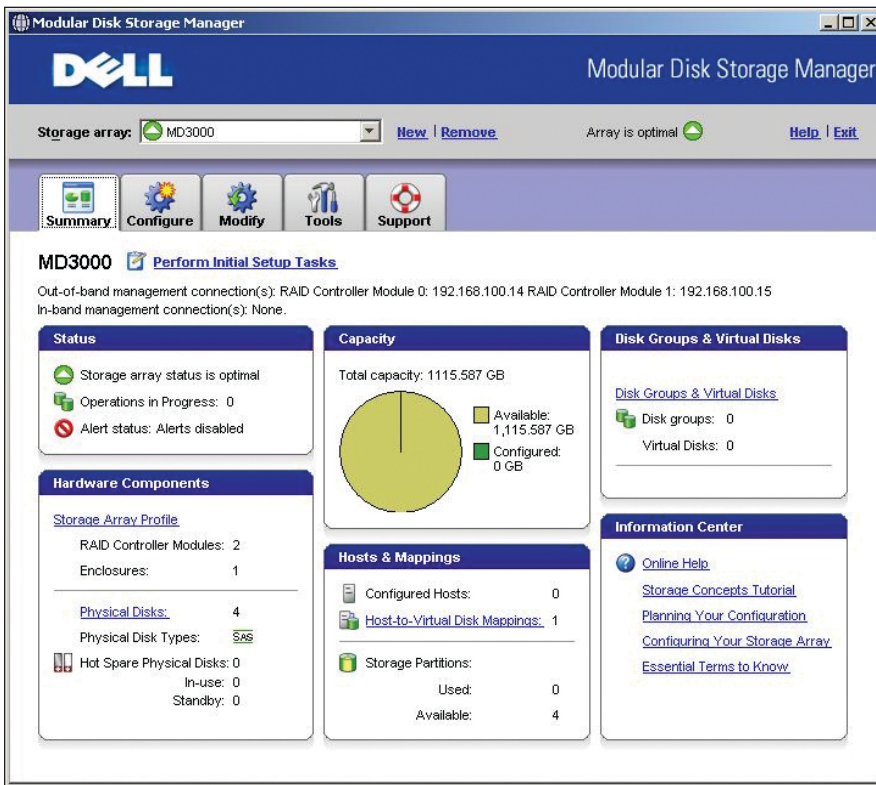


Figure 4. Dell Modular Disk Storage Manager GUI

connected to a host using Ethernet, or a combination of the two

- Display configuration information, add and remove storage arrays, and manage Simple Network Management Protocol (SNMP) and Simple Mail Transfer Protocol (SMTP) alert notifications

Figure 5 shows a script file called `sample.scr` that creates a 2 GB RAID-5 virtual disk called “Test” and sets its parameters. To run this script

file from the CLI, administrators would enter the following command (where `ipaddress` is the IP address of the storage enclosure):

```
smcli ipaddress -f sample.scr;
```

Advanced features of the Dell PowerVault MD3000

Administrators can use the PowerVault MD3000 and Dell Modular Disk Storage Manager to create virtual disk snapshots, virtual disk copies, and

storage partitions. Virtual disk snapshot creation and virtual disk copy creation are premium features that must be purchased and enabled separately.

Virtual disk snapshots

Snapshots are persistent point-in-time images of a specific virtual disk that help eliminate the application downtime traditionally required to create system backups. They typically use less disk space and take less time to create than a physical copy.

Snapshots are implemented using the copy-on-first-write method: when virtual disk data needs to be modified, the original data is copied to the virtual disk snapshot cache before it is modified. The virtual disk on which the snapshot is based—the source virtual disk—must be a standard virtual disk in the storage array. Snapshots are typically created so that software such as a backup application can access them while the source virtual disk remains accessible to hosts.

Modular Disk Storage Manager supports application-aware snapshots, which are recommended for transaction-based applications where all transaction operations must be completed for transactions to be valid. Application-aware snapshots coordinate with the application using the virtual disk to help ensure the consistency of snapshot data with application data. Once the snapshot has been created, the snapshot engine signals the application to resume normal processing. Microsoft Volume Shadow Copy Service (VSS) enables the creation of these snapshots.

Administrators can use the virtual disk snapshot feature to perform the following tasks:

- Create a complete image of data on a virtual disk at a particular point in time
- Provide quick, frequent, nondisruptive backups
- Test new versions of a database system without affecting actual data
- Enable snapshot virtual disks to be read, written, and copied

```
//sample.scr

create virtualDisk diskGroup=1 RAIDLevel=5 userLabel="Test" owner=0
segmentSize=16 capacity=2GB;

set virtualDisk["Test"] mediaScanEnabled=false;
set virtualDisk["Test"] consistencyCheckEnabled=false;
set virtualDisk["Test"] modificationPriority
=high;
```

Figure 5. Sample script for the Dell Modular Disk Storage Manager CLI

- Use virtual disk high-availability features, such as RAID protection and redundant path failover
- Map snapshot virtual disks and make them accessible to any connected host, and make snapshot data available to secondary hosts by mapping the snapshot to the hosts
- Create up to four snapshots per virtual disk
- Increase the capacity of a snapshot virtual disk

Virtual disk copies

Administrators can use the PowerVault MD3000 and Modular Disk Storage Manager to create exact copies of virtual disks. The target virtual disk copy must be at least the same size as the source virtual disk, and write access to the source virtual disk is blocked during the copy operation. Administrators can specify some settings for the virtual disk copy, such as copy priority and read/write permission, before or during the copy operation.

Administrators can use the virtual disk copy feature to perform the following tasks:

- **Copy data to larger disks:** As storage requirements for a virtual disk change, administrators can copy virtual disks to a disk group within the same storage array that uses drives with larger capacity than the current drives.
- **Copy snapshot virtual disks:** Administrators can select snapshot virtual disks as source virtual disks for virtual disk copies, which enables comprehensive backups while helping avoid significant impact on storage array I/O.
- **Assign virtual disk copies to hosts:** Administrators can directly assign hosts to a target virtual disk, although by default the hosts do not have write access to the virtual disks, and an attempt to write to them results in a host I/O error. Administrators

can enable read/write access by changing the target virtual disk permissions from “Read-Only” to “Read/Write.”

- **Recover virtual disks:** If a source virtual disk fails, administrators can un-map this failed virtual disk from its host and map its virtual disk copy to the same host.


Storage partitions

A storage partition is a logical entity consisting of one or more virtual disks that can be accessed by a single host or shared among hosts that are part of a host group. The first time administrators map a virtual disk to a specific host or host group, a storage partition is created; subsequent virtual disk mappings to that host or host group do not create another storage partition. LUN masking is implemented using storage partitions.

Administrators can use single storage partitions or multiple storage partitions for different purposes. Single storage partitions are useful when only one attached host is accessing all of the virtual disks in the storage array, or when all attached hosts share access to all virtual disks on the storage array; this type of configuration requires that all hosts be running the same OS and special software (such as clustering applications) to manage virtual disk sharing and accessibility. Multiple storage partitions are useful when specific hosts are accessing specific virtual disks in the storage array, or when hosts with different operating systems are attached to the same storage array; this type of configuration requires creating a storage partition for each type of host.

Flexible, highly scalable storage

The Dell PowerVault MD3000 is designed to provide a flexible, highly scalable storage system with advanced features such as virtual disk snapshot and virtual disk copy creation.

Deploying this storage system can allow enterprises to implement highly available shared storage and help them avoid the costs and complexities of internal storage. 

Manjusha Gopakumar is a senior engineer analyst in the Test Case Design team at the Dell Bangalore Development Center. She has a bachelor's degree in Computer Science from Rajiv Gandhi Technical University.

Vishal Kadam is a senior engineer analyst in the Test Case Design team at the Dell Bangalore Development Center. He has a diploma in Electronics and Communication and is currently pursuing a degree in Information Technology from the University of Mysore. He is a Microsoft Certified Systems Engineer (MCSE), Sun Certified System Administrator (SCSA), and Brocade Certified Fabric Professional (BCFP).

Paul Winston is a test engineer adviser in the Test Case Design team at the Dell Bangalore Development Center. He has a bachelor's degree in Commerce from Bharathiar University and an advanced diploma in Software Engineering from Aptech. His current interests include implementing storage area network, network attached storage, and cluster solutions.

more
Online
www.dell.com/powersolutions

QUICK LINK

Dell PowerVault MD3000 User's Guide:
support.dell.com/support/edocs/systems/md3000