

Optimizing the Dell PowerVault MD1xxx Family SAS Solution

SAS Solutions Team

Dell | SAS Solutions Team
delltechcenter.com



This white paper is for information purposes only, and may contain typographical errors and technical inaccuracies. The content is provided as is without express or implied warranties of any kind.

Information in this document is subject to change without notice.

© Copyright 2006-2008 Dell Inc. All rights reserved.

Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Dell, the *Dell* Logo, *PowerVault*, *PowerEdge*, and *OpenManage* are trademarks of Dell Inc. *Microsoft* and *Access* are either trademarks or a registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

TABLE OF CONTENTS

1. INTRODUCTION 4
Serial Attached SCSI (SAS)..... 5

2. RECOMMENDED MAINTENANCE FOR DELL POWERVAULT MD1XXX FAMILY SAS SOLUTION 7

3. DELL SAS PRODUCT OFFERING..... 8
SAS Cabling 10
Physical Drive Support 10
Third Party Hardware Support..... 10
Server Support 10

4. POWERVAULT MD1XXX FAMILY STORAGE ENCLOSURE DETAILS..... 11
Enclosure Modes 11
Unified Topology..... 11
Split Topology..... 11
Daisy-Chaining Multiple PowerVault MD1xxx Family Storage Enclosures..... 12
PowerVault MD1xxx Family – Common Features and Capabilities..... 13

5. PERC 5/E AND PERC 6/E CONTROLLER - DETAILED FEATURES 17
PERC 5/E..... 17
PERC 6/E..... 17
Considerations when implementing a RAID Solution 17
Controller Features 17
Capacity Requirements..... 19

6. DELL OPENMANAGE SERVER ADMINISTRATOR..... 21
Managing the PERC 5/E and PERC 6/E Controllers and the PowerVault MD1xxx Family Storage Enclosures..... 21
Using the Create Virtual Disk Wizards 21
Reconfiguring Virtual Disks -- RAID Migration and Capacity Expansion..... 22
Moving Physical Disks and Virtual Disk Configurations – Disk Roaming and Disk Migration 23
Importing Migrated Virtual Disks 24

7. STORAGE APPLICATIONS AND COMPONENTS..... 27
Identify customer usage model and needs..... 27
Picking a solution based on the application profile..... 27
Application-Specific Guidelines..... 30
E-Mail Servers 30
Web Servers 31
Database or Online Transaction Processing Servers..... 32
File Servers..... 34
Streaming Media Servers 37
Hot Spares..... 38

8. SUMMARY AND CONCLUSIONS 39

APPENDIX – A: REFERENCES 40

9. APPENDIX – B: GLOSSARY 41

10. APPENDIX – C: REVISION HISTORY 42

1. Introduction

A complete Enterprise storage solution involves various software and hardware components. A given storage project can be addressed using several approaches. Therefore, it is essential to identify optimum storage solutions for different applications.

Selecting the right storage components not only impacts the effectiveness of a given application but may also improve the performance of the entire IT infrastructure that relies on that storage. See Table 1 for examples of key storage components.

Table 1: Examples of the components in a typical storage solution

Storage Solution Components	Examples
Physical Storage Device Technology	Storage Expansion Enclosure (for example, PowerVault™ MD1xxx Family and MD3xxx Family) ¹
Interconnect Technology	3 Gbps Serial Attached SCSI (SAS)
Host System Interface Technology	Non-RAID (HBA) or RAID controllers (for example, PowerEdge™ Expandable Raid Controller)
Management Capabilities	Dell™ OpenManage™ Server Administrator

As with many IT solutions, storage deployments can be assessed using numerous criteria, each with varying levels of importance. In general, the effectiveness of a given storage solution should be measured using the following attributes:

- *Reliability* – The solution remains functional and ensures that data is unaltered
- *Availability* – Data remains accessible by the application.
- *Serviceability* – The hardware has easy access to assemblies/components, allowing easy maintenance and replacement.
- *Redundancy* – The hardware provides redundancy for both the data storage, as well as for the hardware management capabilities.
- *Performance*
 - Bandwidth Utilization – Data throughput / available Bandwidth
 - IOPS – I/Os per second
- *Data Protection* – Data is protected against catastrophic system failure or components failure.
- *Scalability* – Storage capacity can be expanded as the application necessitates.
- *Cost* – Measured in total lifetime cost, acquisition cost, or \$ per Gigabyte.
- *Density* – The storage solution provides optimal IOPS performance within a given rack space
- *Power and Cooling* – Optimization of power and cooling requirements in the datacenter

¹ Dell PowerVault MD1xxx Family include the MD1000 and MD1120 enclosures. The Dell PowerVault MD3xxx Family include the MD3000 and MD3000i storage arrays.

An optimal storage solution addresses all key attributes to a greater or lesser extent depending on their relative importance. This paper presents a study of the technological, functional and behavioral differences of the storage solution components and measures them against the attributes listed above. Understanding each component with respect to these attributes will help in making informed decisions about the optimal solution for a given storage application.

The paper presumes that the reader has a basic understanding of the different RAID levels and interconnect technologies – such as Serial Attached SCSI (SAS) and PCI Express. This paper focuses on Serial Attached SCSI storage technology, specifically for the Dell PowerVault MD1xxx Family (MD1000 and MD1120) storage expansion enclosures and Dell SAS RAID controllers (PERC 5/E and PERC 6/E). Dell external RAID enclosures (MD3xxx Family) are outside the scope of this document and are discussed under a separate paper.

Other storage technologies such as SCSI or Fibre Channel, as well as topologies such as Network Attached Storage (NAS) or Storage Area Network (SAN), are also outside the scope of this document. All storage applications will be treated independent of a specific host system or server.

Serial Attached SCSI (SAS)


The SCSI (Small Computing System Interface) protocol originated more than 20 years ago and is used in the majority of server storage needs (for example, storage that is internal to a server or physically connected to a server in an external disk expansion enclosure). The market prevalence of SCSI has created economies of scale making it one of the most affordable storage interconnect technologies available today. Many businesses rely on SCSI physical disks to deliver extremely cost-effective and reliable storage. In fact, SCSI physical disks have advanced over the years offering increased performance and more sophisticated features with each iteration. However, given the inherent limitations of its parallel architecture, the current version of SCSI physical disks, U320 or Ultra320, will be the last version of the traditional SCSI technology.

Serial Attached SCSI (SAS) is a major iteration of SCSI technology and introduces a host of benefits when compared to older SCSI technology including the following features.

Dramatic improvement in bandwidth performance


- The first generation of SAS will support 3.0 Gbps (or 300 MB/sec) per “lane”
- Ultra320 SCSI is the shared-bus architecture. SAS introduces a point-to-point topology enabling systems with many physical disks to scale bandwidth far beyond the bandwidth capacity of Ultra320 SCSI.

Figure 1: Parallel U320 SCSI bus versus point-to-point Serial Attached SCSI connection



- SAS introduces the concept of port aggregation to storage interconnects. External storage devices can connect via a x4 (by four) wide-lane. Cabling aggregates four SAS lanes, each operating at 3.0 Gbps, for a full external connection bandwidth of up to 12.0 Gbps on a single connector.

Figure 2: Bandwidth improvement with SAS port aggregation



Improved capacity scalability

- SAS allows for port expansion, enabling the ability to “daisy-chain” storage enclosures for dramatic improvements in capacity scalability. With Ultra320 SCSI, systems were limited to a maximum of 16 devices on a single SCSI channel. This typically translated to maximum of 14 physical disks connected to a single SCSI channel on a RAID controller. With SAS expansion, a single SAS port can be expanded out to multiple SAS devices (similar to how Ethernet hubs can expand a single Ethernet connection into multiple devices).

Other important changes to keep in mind about SAS:

- SAS continues to use the SCSI command sets for driver & software compatibility.
- SAS is not backwards compatible with Parallel SCSI.

2. Recommended Maintenance for Dell PowerVault MD1xxx Family SAS Solution

Dell strives to provide the best possible customer experience for all of our Enterprise server and storage products. Proper maintenance of your solution is essential to ensure that the equipment remains in optimal working condition. Dell recommends the following suggested maintenance activities.

Keep firmware/drivers/utilities Updated

In general, Dell posts planned maintenance updates for firmware twice a year. To stay current with the latest updates, customers can subscribe to Dell notification tools. Click the following link to sign up for technical updates.

<http://support.dell.com/support/topics/global.aspx/support/notifications/en/index?c=us&l=en&s=gen>

Enable Background Read Patrol

Background Read Patrol, a function of the RAID controller, is an automatically scheduled background task that scans and resolves possible physical disk errors that may lead to physical disk failure and result in data integrity issues. Dell recommends implementing Background Read Patrol in all configurations.

See section 5 for details.

Schedule and run periodic consistency checks / Background Initialization (BGI)

Consistency checks / BGI are performed by many RAID controllers to verify that all stripes in a redundant virtual disk are consistent and to automatically correct stripes where inconsistencies are found. BGI is a consistency check that automatically runs on all redundant arrays. Depending on workload, Dell recommends that consistency checks should be run manually on the virtual disks at least once every month to ensure data integrity. Consistency checks and Background Read Patrol are complementary and both should be used.

See section 5 for details.

Enable Hot Spares

Hot spare functionality provides extra security, availability and redundancy by automatically replacing a failed disk in a RAID group and allowing the rebuild of the degraded array to begin immediately. Dell PERC controllers support Hot Spare functionality and it is recommended whenever possible.

See section 5 for details.

Document your configuration and keep the documentation current

Documenting your configuration facilitates faster and easier recovery should a failure occur. Make sure to include RAID levels, physical disks included in your virtual disks, firmware, and driver versions and usage of hot spares in your documentation.

Perform periodic physical inspection of your connections and cabling

Damaged cables and loose connections may cause functional issues with your storage solution. Dell recommends a physical equipment audit two times a year or anytime after re-cabling or re-deploying your storage hardware.

3. Dell SAS Product Offering

The PowerVault MD1xxx Family enclosures represent Dell JBOD offering. This includes both the MD1000 and MD1120 enclosures (see Table 2). These are feature-optimized storage expansion enclosures that use the latest industry standard storage interconnect technology, Serial Attached SCSI (SAS). These enclosures support between 2 and 24 physical SAS drives. Along with Dell host-based RAID controllers, PERC 5/E and PERC 6/E, the PowerVault MD1xxx Family introduces new levels of performance, capacity, and scalability at industry leading prices.

Table 2: Summary and comparison of key features of the MD1xxx Family enclosures

Features	PowerVault MD1000	PowerVault MD1120
Rack space	3U (rack only)	2U (rack only)
Drive bays	15 (3.5-inch drives)	24 (2.5-inch drives)
Supported drive types	3 Gbps SAS & 3 Gbps SATA II	3 Gbps SAS
Expandability	Three shelves for a total of 45 drives per host port (unified mode only)	Three shelves for a total of 72 drives per host port (unified mode only)
Management	In-band, using Dell OpenManage software	In-band, using Dell OpenManage software
Backplane options	Unified or split JBOD	Unified or split JBOD
Supported RAID controllers	PERC 5/E and PERC 6/E	PERC 6/E only
Cluster support	No	No
Hot-pluggable drives	Yes	Yes
Hot-pluggable fans and power supplies	Yes (combined fan and power supply modules)	Yes (combined fan and power supply modules)
Enclosure management configurations	Redundant and nonredundant	Redundant and nonredundant
Fan and power supply configuration	Redundant	Redundant
Supported disk array enclosures	PowerVault MD3000 and PowerVault MD3000i	N/A (designed for future support for attachment to external RAID enclosures)

Primary use	General-purpose storage utilizing cost-effective, high-capacity drives	High-performance storage for applications requiring high I/O or throughput rates
--------------------	--	--

Major features of the PowerVault MD1xxx Family and PERC 5/E and PERC 6/E solution include:

Performance

- Support for 3.0 Gbps SAS physical disks (both 10,000 and 15,000 RPM)
- External cable supports a SAS x4 wide-lane to achieve a total interface bandwidth of up to 12.0 Gbps connection to the host controller
- Host interconnect with the PERC 5/E and PERC 6/E host-based controllers
- RAID levels 0 and 10
- Refer to following link for specific performance information
<http://www.delltechcenter.com/page/Storage>

Availability

- Support for RAID levels 1, 5, 6, 10, 50, and 60
- Redundant, hot pluggable power supplies with integrated cooling fan modules
- Redundant enclosure management modules (EMMs)
- Hot pluggable physical disks

Scalability/Expandability

- Support for daisy-chaining up to three PowerVault MD1xxx Family expansion enclosures per PERC 5/E and PERC 6/E connection
- Supports between 2 and 24 physical drives
- Supports 36GB, 73GB, 146GB, 300, and 400 GB 3.5-inch SAS physical disks (MD1000)
- Supports 250GB, 500GB, 750GB, and 1TB 3.5-inch SATA physical disks (MD1000)
- Supports 36GB, 73GB, 146GB GB 2.5-inch SAS physical disks (MD1120)
- Refer to the Dell PowerVault MD1120 System Support Matrix on www.dell.com for the latest updates

Ease-of-Use

- Common management software with PowerEdge™ servers through Dell OpenManage Server Administrator

PERC 5/E and PERC 6/E are the RAID controllers used to connect the PowerVault MD1xxx Family storage enclosures to one or two PowerEdge servers. They support 3 Gbps Serial Attached SCSI (SAS) as the storage interconnect technology and PCI Express (PCI-E) as the host-based interconnect technology. The “E” in PERC 5/E and PERC 6/E stands for “**External**” because the PERC controllers supports only physical disks that attach to a server externally.

The PERC 5/E controller offers:

- Intel IOP333 I/O processor based on Intel XScale Microarchitecture
- 256MB of customized DDR2 400MHz, Error-Correcting Code (ECC) cache memory
- 3 Gbps maximum speed for each SAS lane
- Two external x4 (by four) SAS ports, each aggregating 4 SAS lanes for a total bandwidth per port of 12.0 Gbps
- x8 PCI E host interface for a total bandwidth of 32.0 Gbps
- Up to 72 hours of intelligent, transportable, battery-backed, cache memory

The PERC 6/E controller offers:

- LSI 1078 SAS RAID-on-Chip
- 256MB or 512MB of customized DDR2 667MHz, Error-Correcting Code (ECC) cache memory
- 3 Gbps maximum speed for each SAS lane
- Two external x4 (“by four”) SAS ports, each aggregating 4 SAS lanes for a total bandwidth per port of 12.0 Gbps
- x8 PCI E host interface for a total bandwidth of 32.0 Gbps
- Up to 72 hours of intelligent, transportable, battery-backed, cache memory (24-hour back-up for 512MB memory)

For detailed information on the PERC 5/E and PERC 6/E controllers, see section 6.

SAS Cabling

The PowerVault MD1xxx Family storage enclosures and the PERC 5/E and PERC 6/E controllers support industry standard 3.0 Gbps SAS cables in 0.6m (meter), 1m, 2m, and 4m cable lengths. Each cable supports x4 SAS connectors to deliver up to 12.0 Gbps of total bandwidth from the host to the enclosure. Only Dell-qualified cables are supported for connecting the PowerVault MD1xxx Family storage enclosures to the PERC 5/E and PERC 6/E controllers.

Physical Drive Support

The PowerVault MD1xxx Family storage expansion enclosures supports the latest SAS physical disks, currently specified as 3.0 Gbps SAS. SAS solutions from Dell include both 10,000 RPM and 15,000 RPM spindle speeds and a variety of capacities to allow customers the maximum flexibility.

The PowerVault MD1000 enclosure also supports 7200 RPM SATA drives, allowing for higher capacity than SAS drives.

For the latest information on support for specific physical disks, please refer to the respective PowerVault MD1000 and MD1120 System Support Matrices at the following links:

For MD1000 - <http://support.dell.com/support/edocs/systems/md1000/>

For MD1120 - <http://support.dell.com/support/edocs/systems/md1120/>

Third Party Hardware Support

To provide the best possible customer experience, Dell supports only hardware purchased direct from Dell for the PowerVault MD1xxx Family and PERC 5/E and PERC 6/E controller. Dell developed server and storage solutions include many thousands person-hours of engineering testing and validation to ensure the highest product quality. For example, physical disks purchased direct from Dell may include features not generally available to the rest of the industry. As such, Dell supports only Dell tested and qualified hardware. Hardware purchased direct from Dell carries the full Dell system warranty and ensures full compatibility. Materials purchased from third-party vendors for use in Dell systems are not covered under Dell warranty and may not deliver an acceptable customer experience.

Server Support

The PowerVault MD1xxx Family is supported on an broad line of Dell server systems. For the latest information on support for specific servers, please refer to the respective PowerVault MD1000 and MD1120 System Support Matrices at the following links:

For MD1000 - <http://support.dell.com/support/edocs/systems/md1000/>

For MD1120 - <http://support.dell.com/support/edocs/systems/md1120/>

4. PowerVault MD1xxx Family Storage Enclosure Details

The Power Vault MD1xxx Family includes two JBOD enclosures, the MD1000 and the MD1120. These are feature-optimized storage expansion enclosures which utilize the latest industry standard storage interconnect technology, Serial Attached SCSI (SAS). These enclosures support between 2 and 24 physical SAS drives. This section outlines supported enclosure mode features and configuration topologies that apply to both the MD1000 and MD1120 enclosures.

Note: MD1000 is supported on both the PERC 5/E and PERC 6/E controllers. MD1120 is only supported on the PERC 6/E controller.

Enclosure Modes

The enclosures can be set to either unified or split mode through the use of a two-position switch located on the front panel of the enclosures. An enclosure can be configured into the selected mode only at power up. Any changes to the mode switch made while the enclosure is powered on will not take effect until the enclosure is powered off and back on again.

Note: Any time you power on a PowerVault MD1xxx Family storage enclosure, ensure the mode switch is set in the desired position. If the position of the switch changes inadvertently, data loss could occur as the controller may recognize an inaccurate number of attached physical disks versus the previous mode. In order to cycle power on your MD1xxx Family storage enclosure, all attached servers must be turned off.

Unified Topology

In unified mode, a single server has access to all physical disks in a PowerVault MD1xxx Family storage enclosure through a single SAS cable. The unified mode can also be used in a daisy-chaining configuration to provide a single server access to up to three enclosures from a single host port. The main advantage of unified mode is that it provides a high degree of physical disk scalability.

Split Topology

In split mode, an enclosure is divided into two logical enclosures; with each host connection accessing its own set of physical disks. The division between EMM ownership of physical disks is indicated by a disk separation icon in the same area as the physical disk numbering labeling on the front of the enclosure chassis. See the table below for physical disk ownership for specific MD1xxx Family enclosures. To connect a server to an enclosure in split mode, the **In** port of each EMM should be attached to a connector of a PERC 5/E controller. The **Out** port of an EMM in split mode is disabled, since enclosure daisy-chaining is not supported in this mode. See Figure 3 for an example of an enclosure in a split mode configuration. An enclosure in split mode can also have both **in** ports connected to different channels of a single host controller.

EMM Physical Disk Ownership in Split Topology

MD1xxx Family Storage Enclosure	Primary EMM Physical Disks	Secondary EMM Physical Disks
MD1000	7 through 14	0 through 6
MD1120	12 through 23	0 through 11

Split mode is useful in deployments where a single PowerVault MD1xxx Family storage enclosure provides limited amounts of storage each for two different servers.


Note: In split mode, each host server can see its own physical disks but not the other server's physical disks.

Note: Split mode does not support daisy-chaining.

Note: A system with a single EMM set in split mode will function properly, but the server will only be able to see the physical disks that are connected to the EMM.

Dual Hosts

Figure 3: Split mode configuration



Daisy-Chaining Multiple PowerVault MD1xxx Family Storage Enclosures

In unified mode, up to three enclosures may be daisy-chained together from a single PERC 5/E or PERC 6/E controller port. A maximum of six PowerVault MD1xxx Family storage enclosures (3 enclosures per port) may be connected to a single PERC 5/E or PERC 6/E controller. (Note: Although this is a supported configuration, one should consider total solution availability before attaching this quantity of storage to a single-controller solution.). When connecting multiple enclosures, the first enclosure is connected to the PERC 5/E or PERC 6/E controller via the **In** port of the primary EMM. The second enclosure is connected from the **Out** port of the primary EMM on the first enclosure to the **In** port of the primary EMM on the second enclosure. The connection is made with a standard SAS cable. The third enclosure is connected in the same manner as the second enclosure.

Note: Mixing MD1xxx Family storage enclosure types in a daisy-chain is not supported.

Note: Mixing enclosure types on different channels of a single host controller is not supported. Only one MD1xxx Family storage enclosure type is supported on a host controller at one time. Two host controllers can support two enclosure types: one enclosure type on each controller.

Figure 4: Max configuration on a single PERC 5/E or PERC 6/E controller port





Figure 5: Max configuration on a single PERC 5/E or PERC 6/E controller (both ports connected)



PowerVault MD1xxx Family – Common Features and Capabilities

Regardless of the topology, the PowerVault MD1xxx Family storage enclosure provides a set of common features and capabilities.

World Wide Name (WWN) Assignment

The SAS standard requires that all devices on a SAS domain have a unique address, called a World Wide Name or WWN. PowerVault MD1xxx Family storage enclosures comply with the SAS standard. For more information on WWN assignment, please refer to the SAS Specification referenced in Appendix A.

Mixing Physical Disk Capacities

By supporting the mixing of physical disk of different sizes within an enclosure, PowerVault MD1xxx Family storage enclosures enable optimal matching of physical disk type to the requirements of multiple applications. Mixing of physical disk sizes is supported but not

recommended within a single RAID group. PowerVault MD1000 supports 36 GB, 73 GB, 146 GB, 300 GB, and 400 GB capacities for 3.5" SAS physical disks. PowerVault MD1120 supports 36 GB, 73 GB, and 146 GB 2.5" SAS physical disks. When operating in a RAID configuration, the lowest capacity physical disk defines the RAID virtual disk size. For example, a RAID 1 configuration with two different sized physical disks, one 36 GB and another 146 GB physical disk, will be limited to 36 GB. If the same two physical disks were configured in RAID 0 configuration, the virtual disk size would be limited to 73 GB (the equivalent of two 36 GB physical disks in RAID 0 configuration).

For the latest information on support for specific physical disks, see the Dell Storage web site at:

<http://www.us.dell.com/storage/>

Mixing Physical Disk Spindle Speeds

Mixing different physical disk spindle speeds is supported, but not recommended within a single RAID group. PowerVault MD1xxx Family storage enclosures support 10,000 and 15,000 RPM spindle speeds. Depending upon the RAID configuration, the lowest spindle speed may dictate the maximum performance of the RAID group.

SCSI Enclosure Services (SES)

The EMM continuously monitors environmental variables such as temperature and voltage. PowerVault MD1xxx Family storage enclosures work in conjunction with Dell OpenManage Server Administrator to set trigger thresholds that warn the user when these variables exceed safe operating levels. The enclosure will automatically power down when these environmental variables exceed a critical threshold level in order to avoid potential equipment damage. Both local visual indicators (for example, LEDs) and audio alerts are triggered when faults are detected. The audio alerts on the enclosure are disabled by default, but can be enabled using Server Administrator. The following SES functions are provided by the EMMs in the PowerVault MD1xxx Family storage enclosures:

- Enclosure shelf faults detection
- Physical disk Power/Fault/Offline status
- Power supply status and shutdown capabilities
- Fan speed status and control
- Temperature monitoring
- Over/under temperature shutdown capabilities
- Alarms (local audible and visual LED indicators)

Hot Plugging

- PowerVault MD1xxx Family storage enclosures are designed for hot plugging certain components to prevent damage when pulling or replacing. Hot plugging is a concept which allows for the safe removal and replacement of components while the power is still on. Hot plugging allows server applications and disk activities to continue uninterrupted, while maintenance or repair actions take place. Additional steps may be required for physical disks to become usable. See the RAID controller User's Guide for more information.
- Hot Pluggable Devices
 - Physical disk (all Disk I/O activities must be halted first)
 - Power supply module/ cooling fan module

Note: Hot plugging of SAS cables is not supported.

External SAS Cables

Depending on the enclosure topology, one or two external SAS cables may be attached between the host system(s) and a PowerVault MD1xxx Family storage enclosure. SAS cables use four active 3 Gbps SAS lanes in a single cable to deliver an aggregate of 12.0 Gbps connection. All PowerVault MD1xxx Family SAS cables purchased from Dell have been tested and comply with the ANSI SAS 1.0 standard. Dell supports using only Dell-provided cables.

External SAS Connector Types

SAS uses a 25-pin connector whose specification is governed by the SFF-8470 specification. For more information, refer to the specification referenced in Appendix A.

Regardless of the physical topology, PowerVault MD1xxx Family storage enclosures also provide some common capabilities as listed below:

Redundancy Capabilities

- *Management Redundancy*

With two EMMs installed, a PowerVault MD1xxx Family storage enclosure's internal enclosure functions such as temperature regulation, LED control, and alarm control, are maintained in the event of an EMM failure. However, even with dual EMMs installed, enclosure status will not be reported to the host in the event of an EMM failure.

- *Power Supply and Cooling Redundancy*

PowerVault MD1xxx Family storage enclosures require that both power supplies be installed for normal operation. A system can operate using a single power supply for a limited time in the event a power supply fails or is removed for service.

PowerVault MD1xxx Family storage enclosures provide individually controlled fans. The units is thermally designed to continue operating in the event of one fan blower failure (n+1 cooling redundancy). In the event of a fan failure, the power supply containing the failed fan must be replaced.

Reliability Capabilities

- *Thermal Reliability*

A fully-redundant PowerVault MD1xxx Family storage enclosure contains four digital temperature sensors, two on the backplane and one on each EMM. In conjunction with the SES processor, the temperature sensors provide device protection against adverse thermal conditions. PowerVault MD1xxx Family storage enclosures are designed to operate within ambient temperature range of 10°C (50°F) to 35°C (95°F). When either of the backplane sensors detects temperature below 3°C (37.4°F) or above 55°C (131°F), the SES processor notifies the host of a critical temperature situation. This triggers the host to perform a graceful shutdown of the enclosure within three minutes. In the event that the host is unable to initiate the enclosure shutdown process or a host is not connected to the enclosure, a PowerVault MD1xxx Family storage enclosure has an

embedded thermal shutdown feature. This is activated by the SES processor in either of the following cases:

- Either of the backplane temperature sensors detects temperature below -2°C (28.4°F) or above 60°C (140°F).
- Temperature sensor on an EMM detects temperatures below -2°C (28.4°F) or above 65°C (149°F).

A PowerVault MD1xxx Family storage enclosure shuts down within 5 seconds of the activation of an embedded thermal shutdown.

PowerVault MD1xxx Family Serviceability Capabilities

- All the major PowerVault MD1xxx Family components are easily accessible and replaceable even when in a rack.
- PowerVault MD1xxx Family storage enclosures are equipped with LED indicators for visual status reporting on the status of internal components.
- A PowerVault MD1xxx Family storage enclosure's audible alarm notifies users of critical/non-critical status. (For details, see PowerVault MD1xxx Family *Hardware Owners Manual*). Status is also reported to the host system's management tool, Server Administrator. (For details, see the OpenManage User's Guide.)
- PowerVault MD1xxx Family storage enclosures use the Dell Update Utility (DUP) for downloading firmware during serviceability. In a daisy-chain configuration, all EMMs in the configuration are updated via a single DUP instance on the host. After a firmware download, all updated enclosures and the host server must be power cycled for the changes to take effect.

5. PERC 5/E and PERC 6/E Controller - Detailed Features

PERC 5/E

PERC 5/E is the RAID controller used to connect the PowerVault MD1000 storage enclosure with PowerEdge servers. It supports 3 Gbps Serial Attached SCSI (SAS) as the storage interconnect technology and PCI -E as the host-based interconnect technology. PERC 5/E controller supports RAID levels 0, 1, 5, 10 and 50.

PERC 6/E

PERC 6/E is the RAID controller used to connect the PowerVault MD1xxx Family of storage enclosures with PowerEdge servers. It supports 3 Gbps Serial Attached SCSI (SAS) as the storage interconnect technology and PCI -E as the host-based interconnect technology. PERC 6/E controller supports RAID levels 0, 1, 5, 6, 10, 50, and 60.

Considerations when implementing a RAID Solution

The selection of a RAID solution in any environment is governed by several aspects including controller and capacity requirements. Each of these factors is discussed in detail in the following sections.

Controller Features

The Server Administrator is the recommended application for configuring and managing the PowerVault MD1xxx Family. Apart from the Server Administrator software, the common BIOS configuration utility (<CTRL> <R>) of PERC 5/E and PERC 6/E RAID controllers may be used to perform pre-OS configuration and fault recovery. Some of the key features of the PERC 5/E and PERC 6/E controller are listed here:

- *Virtual Disk Initialization*

PERC 5/E and PERC 6/E controllers are capable of performing initialization on the virtual disk(s) by either writing 0's to a portion of the physical disk (fast initialization) or to the entire physical disk (full initialization). Full initialization guarantees that any legacy data is cleared and all sectors are healthy.

- *Background Patrol Read*

Background Patrol Read is an automatically scheduled background task that scans and resolves possible physical disk errors that may lead to physical disk failure and result in data integrity issues. Dell recommends implementing Background Patrol Read in all the configurations.

- *Consistency Check/Background Initialization (BGI)*

Consistency checks / BGI are performed by PERC 5/E and PERC 6/E controllers to verify that all stripes in a redundant virtual disk are consistent and to automatically correct stripes where inconsistencies are found. BGI is a consistency check that automatically runs on all redundant arrays. Depending on workload, Dell recommends that consistency checks are run manually on virtual disks at least once every month to ensure data integrity. Running a consistency check is complement to running Patrol Read and both should be used.

- *RAID Level Migration*

Changing the RAID level of a virtual disk from the existing level (for example, RAID 1) to another level (for example, RAID 5) is called RAID level migration. Dell recommends that data in the virtual disk be backed up before a RAID level migration is initiated. Dell also strongly recommends that the size of the physical disks to which the RAID level is to be migrated is greater than or equal to the size of the physical disks in the existing virtual disk. Dell supports RAID level migration as follows:

RAID 0 to RAID 1
RAID 0 to RAID 5
RAID 0 to RAID 6*
RAID 1 to RAID 0
RAID 1 to RAID 5
RAID 1 to RAID 6*
RAID 5 to RAID 0
RAID 5 to RAID 6*
RAID 6 to RAID 0*
RAID 6 to RAID 5

* RLM only supported on PERC 6/E only.

- *Capacity Expansion*

Expanding the capacity of a virtual disk by adding another physical disk is called capacity expansion. Dell supports capacity expansion on all the supported RAID levels except RAID 1. Dell recommends that the data in the virtual disk be backed up before capacity expansion is initiated. For more details on capacity expansion, see the user's guide for your specific controller.

- *Controller Write Cache Policy Options*

Write-Back Caching

In write-back caching, data transfer is completed when the controller cache receives all data from host for the write transaction. Write-back caching is faster than write through caching. The main advantage of write-back caching is improved performance since the controller does not have to wait for an acknowledgement from the physical disk before proceeding. The battery backup unit for the controller is required to enable write-back caching. For more details on write-back caching, see the user's guide for your specific controller.

Write-Through Caching

In write-through caching, a data transfer is completed when the disk subsystem receives all of the data from the host. Write-through caching provides additional security because data must be committed to the physical disk before proceeding. There may be a minimal performance impact since the controller must wait for the physical disk to return a good status to the controller before proceeding to the next operation.

- *Hot Spare Support*

When a physical disk is assigned to a virtual disk as a hot spare, the PERC 5/E and PERC 6/E controllers will automatically rebuild the virtual disk using the hot spare in the event of a failed physical disk. This restores the virtual disk to optimal redundancy. The hot spare disk must be equal to or greater than the largest physical disk size within any of the virtual disks. Dell recommends assigning hot spares to redundant virtual disks only, (i.e. RAID 1, RAID 5, RAID 6, RAID 10, RAID 50 and RAID 60.)

Global Hot Spare

A global hot spare is a backup physical disk capable of replacing a failed physical disk in any virtual disk on PERC 5/E and PERC 6/E controllers. The capacity of a global hot spare should be greater than or equal to the largest physical disk in all the disk groups.

Dedicated Hot Spare

A dedicated hot spare is a backup physical disk capable of replacing a failed physical disk in a specific PERC 5/E or PERC 6/E virtual disk. The capacity of a dedicated hot spare should be greater than or equal to the largest physical disk in the virtual disk to which it is assigned. Dell recommends assigning dedicated hot spares to all business critical virtual disks.

- *Disk Roaming*

Disk roaming is initiated when the physical disks are changed to different ports on the same controller. The physical disk roaming feature will detect the RAID configuration from the configuration metadata residing on the physical disk.

- *Rebuild*

The process of restoring the redundancy of a RAID 1, 5, 6, 10, 50 or 60 virtual disk is called rebuilding. The rebuilding process can be initiated manually or automatically. Dell recommends choosing the automatic rebuild option to prevent running in the degraded mode for an extended period of time.

- *Mixed-Size Physical Disk Within RAID 10, RAID 50 or RAID 60*

When mixing physical disk sizes in RAID 10, RAID 50 or RAID 60, the controller will not coerce the larger virtual disks to match the size of the smallest virtual disk. PERC 5/E and PERC 6/E controllers will stripe across all disks until the smallest one is full, then it will continue to stripe across the remaining disks until the next one is full, and so on until all the available space is used. This process can impact performance as the array fills up, since the numbers of stripes decreases as disks fill up.

Capacity Requirements

The PERC 5/E and PERC 6/E controllers support up to sixty four virtual disks simultaneously of which each virtual disk can contain up to thirty two physical disks. The maximum number of total virtual disks that can be supported on each physical disk is sixteen. PERC 5 and PERC 6/E controllers support virtual disks greater than 2 TB in size, but Dell recommends checking the capacity of the operating systems to support boot volumes of that capacity.

<http://www.microsoft.com/windowsserversystem/storage/getstorfacts.mspx>

http://www.dell.com/downloads/global/vectors/2004_2tblun.pdf

Table 3: PERC 5/E and PERC 6/E controller features

Feature	PERC 5/E	PERC 6/E
SAS Technology	3Gbps	3Gbps
Host Bus Interface	x8 PCI Express	x8 PCI Express
Internal connectors	N/A	N/A
External Connectors	2	2
RAID Levels Supported	0, 1, 5, 10 and 50	0, 1, 5, 6, 10, 50 and 60
Stripe Sizes Supported	8 KB, 16 KB, 32 KB, 64 KB and 128 KB	8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB and 1024 KB,
Cache Size (DIMM)	256MB DDR2	256MB or 512MB DDR2
Write Cache Options	Write-Through, Write-Back	Write-Through, Write-Back
Read Cache Options	Read-Ahead, No Read Ahead	Read-Ahead, No Read Ahead
BBU Battery Life	Up to 72 hours	Up to 72 hours for 256MB Up to 24 hours for 512MB
Maximum Number of Virtual Disks/Controller	64	64
Maximum Number of Physical Drives/Virtual Disk	32	32
Form Factor	Full Height, Half-Length PCI Adapter card	Full Height, Half-Length PCI Adapter card
I/O Processor (IOP)	Intel IOP333	LSI 1078 ROC

6. Dell OpenManage Server Administrator

Dell OpenManage Server Administrator is a management application for configuring PowerEdge servers. The Server Administrator's configuration abilities for host-based RAID and non-RAID disk storage are of particular interest for this paper. The Server Administrator enables you to perform controller and enclosure functions from a single, graphical or command line interface. The graphical user interface (GUI) is wizard-driven with features for novice and advanced users and detailed online help. The command-line interface is fully-featured and scriptable.

The Server Administrator provides SAS support with the PERC 5/E and PERC 6/E controllers and the PowerVault MD1xxx Family storage enclosures. Server Administrator also supports the SCSI and ATA for RAID technologies.

Managing the PERC 5/E and PERC 6/E Controllers and the PowerVault MD1xxx Family Storage Enclosures

After installing Server Administrator, the user can expand the storage object in the tree view to display the storage components attached to the system. These components include supported controllers and enclosures, virtual disks, physical disks, and the connectors (SAS ports and/or SCSI channels) that comprise the topography of all attached storage.

The PERC 5/E and PERC 6/E controllers and the PowerVault MD1xxx Family storage enclosures are represented in Server Administrator's tree view. You can select the PERC 5/E and PERC 6/E objects and the PowerVault MD1xxx Family object to display the status and to provide access to PERC 5/E and PERC 6/E controllers and PowerVault storage enclosure tasks.

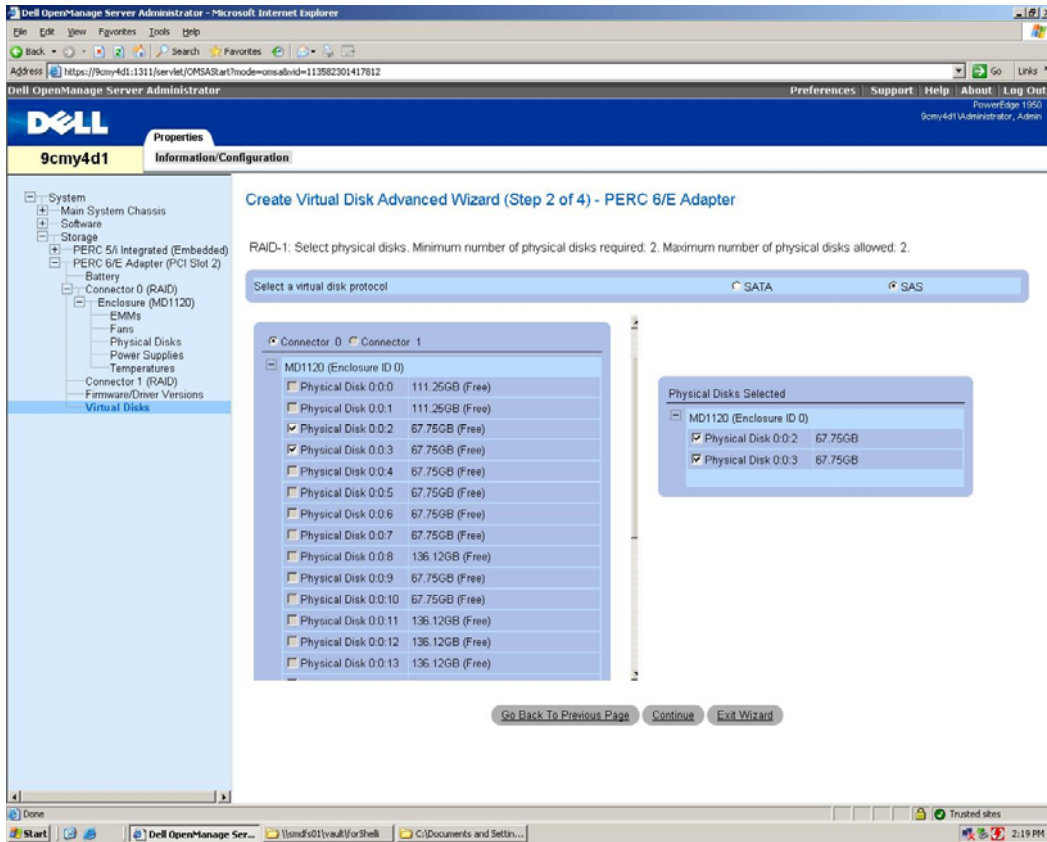
Using the Create Virtual Disk Wizards

You can create virtual disks using either Server Administrator's *Create Virtual Disk Express* or the *Create Virtual Disk Advanced* wizard. These wizards enable you to quickly configure virtual disk attributes, such as the RAID level, size, and physical disk selection.

The *Express* wizard displays the RAID levels supported by the controller and enables you to select the desired RAID level and the disk size for a virtual disk. After you select the RAID level, the *Express* wizard selects an appropriate number of physical disks for the RAID level. If there is available space on disks of this RAID level, these disks will be used for the creation of the virtual disk. Otherwise, new disks of maximum size are used to create the virtual disk. The *Express* wizard displays the physical disks it has selected so that you can confirm the selection before creating the virtual disk. The *Express* wizard also assigns a hotspare to your virtual disk, if applicable.

The *Advanced* wizard allows you to select physical disks and specify additional virtual disk properties such as the read, write, and cache policies.

Figure 6: Screenshot of Create Virtual Disk Advanced Wizard with Disk Selection Displayed



The *Create Virtual Disk Express* and *Create Virtual Disk Advanced* wizards set the virtual disk initialization to “Fast Initialize” by default. After completing the creation of a virtual disk using the Express or Advanced wizard, the PERC 5/E and PERC 6/E controllers implement the virtual disk configuration on the selected physical disks and completes a fast initialize.

Reconfiguring Virtual Disks -- RAID Migration and Capacity Expansion

Server Administrator supports RAID migration and capacity expansion with the *Reconfigure Virtual Disk* wizard. This wizard enables you to change a virtual disk’s RAID level or increase its size by adding one or more physical disks.

The possibilities for reconfiguring a virtual disk depend on the capabilities of the RAID controller, the existing RAID level, and the available physical disks. The following table describes the possible reconfiguration scenarios on the PERC 5/E and PERC 6/E controller.

Table 4: Virtual disk reconfiguration on the PERC 5/E and PERC 6/E RAID controllers

PERC 5/E and PERC6/E Controller Starting RAID Level	Target RAID Level	Disk Requirements
RAID 0	RAID 0	Add at least one additional disk
RAID 0 (on a single disk)	RAID 1	Add a single disk
RAID 0 (on a single disk)	RAID 5	Add at least two additional disks
*RAID 0 (on a single disk)	RAID 6	Add at least three additional disks
RAID 1	RAID 0	With or without adding additional disks
RAID 1	RAID 5	Add at least one additional disk
*RAID 1	RAID 6	Add at least two additional disks
RAID 5	RAID 0	With or without adding additional disks
*RAID 5	RAID 6	Add at least one additional disk
*RAID 6	RAID 0	With or without adding additional disks
*RAID 6	RAID 5	With or without adding additional disks

* denotes only applicable to the PERC 6/E controller

Other considerations may apply to virtual disk reconfiguration. For example, you cannot create more than sixty four virtual disks on the PERC 5/E and PERC 6/E controller. After reaching this limit, you will not be able to reconfigure any of the virtual disks on the controller.

The *Reconfigure Virtual Disk* wizard displays the available target RAID levels and the physical disks. After you have selected the new RAID level and the physical disks, the *Reconfigure Virtual Disk* wizard displays the previous and the new configuration so that you can confirm or reject the changes.

Moving Physical Disks and Virtual Disk Configurations – Disk Roaming and Disk Migration

Disk roaming refers to physically moving a disk from one cable connection to another or from one backplane or enclosure slot to another on the same controller. The PERC 5/E and PERC 6/E controllers recognize the relocated physical disks and logically restores the physical disk and its data to the proper virtual disk. Disk roaming can be performed only when the system is shut down.

Disk migration refers to moving physical disks from one controller to another. Virtual disks residing on the relocated physical disks are identified as a *foreign configuration* on the receiving PERC 5/E or PERC 6/E controller. The Server Administrator enables you to import the foreign configuration so that the migrated virtual disks can be managed on the receiving PERC 5/E or PERC 6/E controller.

The following conditions must be met to successfully migrate the physical disks and their resident virtual disks:

- Virtual disks on a SAS controller can only be migrated to another SAS controller.
- When moving an enclosure, power down the enclosure and the server before moving the enclosure. The enclosure can be moved to any connector number on the receiving controller.

The following procedure describes how to migrate the physical disks from one controller to another:

1. Shut down the system from which the physical disks are being moved.
2. Shut down the server of the receiving controller.
3. Place the physical disks into the new enclosure.
4. Start up the system connected to the receiving controller.
5. Use Server Administrator's *Import/Recover Foreign Configuration* or *Import Foreign Configuration* wizard to import the migrated virtual disks on the receiving controller.

NOTE: If no virtual disks were on the receiving controller, then the foreign physical disks will be auto imported. (I.e. they will not appear in the foreign state)

Importing Migrated Virtual Disks

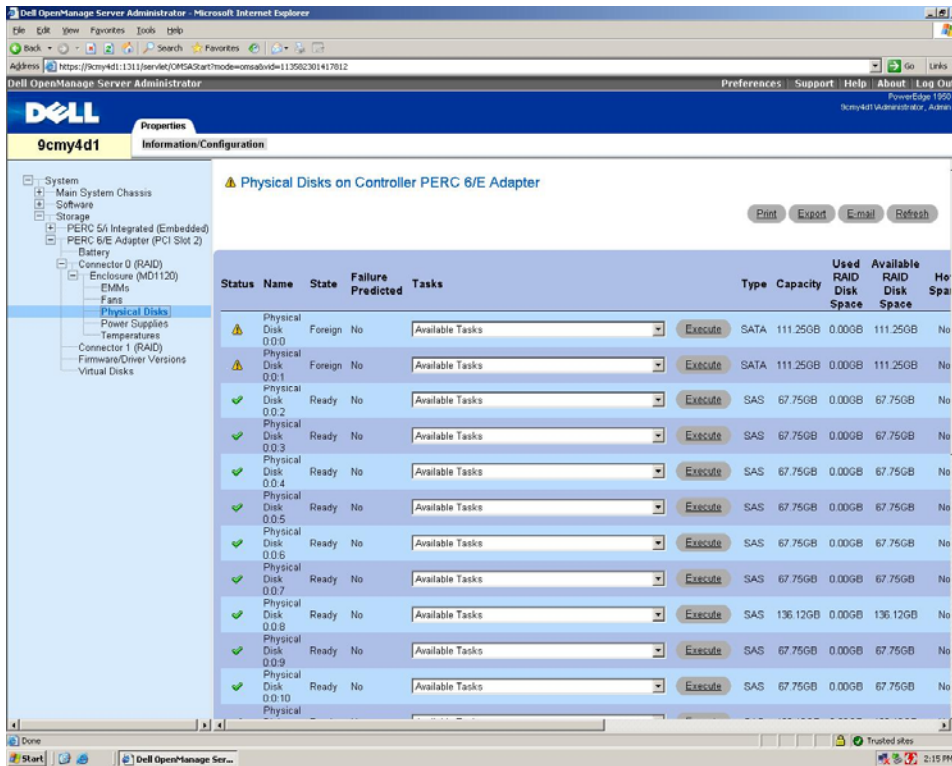
After migrating the physical disks, Open Manage Server Administrator Storage Management enables you to import the virtual disks or "foreign" configuration that resides on the relocated physical disks.

To be imported, the foreign configuration must contain a virtual disk that is in either "Ready" or "Degraded" state. For proper completion of this task, all of the virtual disk data must be present. If the virtual disk is using a redundant RAID level, the additional redundant data does not need to be present. For example, if the foreign configuration contains one side of a mirror in a RAID 1 virtual disk, the virtual disk is in a degraded state and can be imported. On the other hand, if the foreign configuration contains only one physical disk that was originally configured as a RAID 6 using four physical disks, then the RAID 6 virtual disk cannot be imported.

In addition to the virtual disks, a foreign configuration may consist of a physical disk that was assigned as a hot spare on one controller and then moved to another controller. The *Import/Recover Foreign Configuration* or *Import Foreign Configuration* wizard imports the new physical disk as a hot spare. If the physical disk was a dedicated hot spare on the previous controller, but the virtual disk to which the hot spare was assigned is no longer present in the foreign configuration, then the physical disk is imported as a global hot spare.

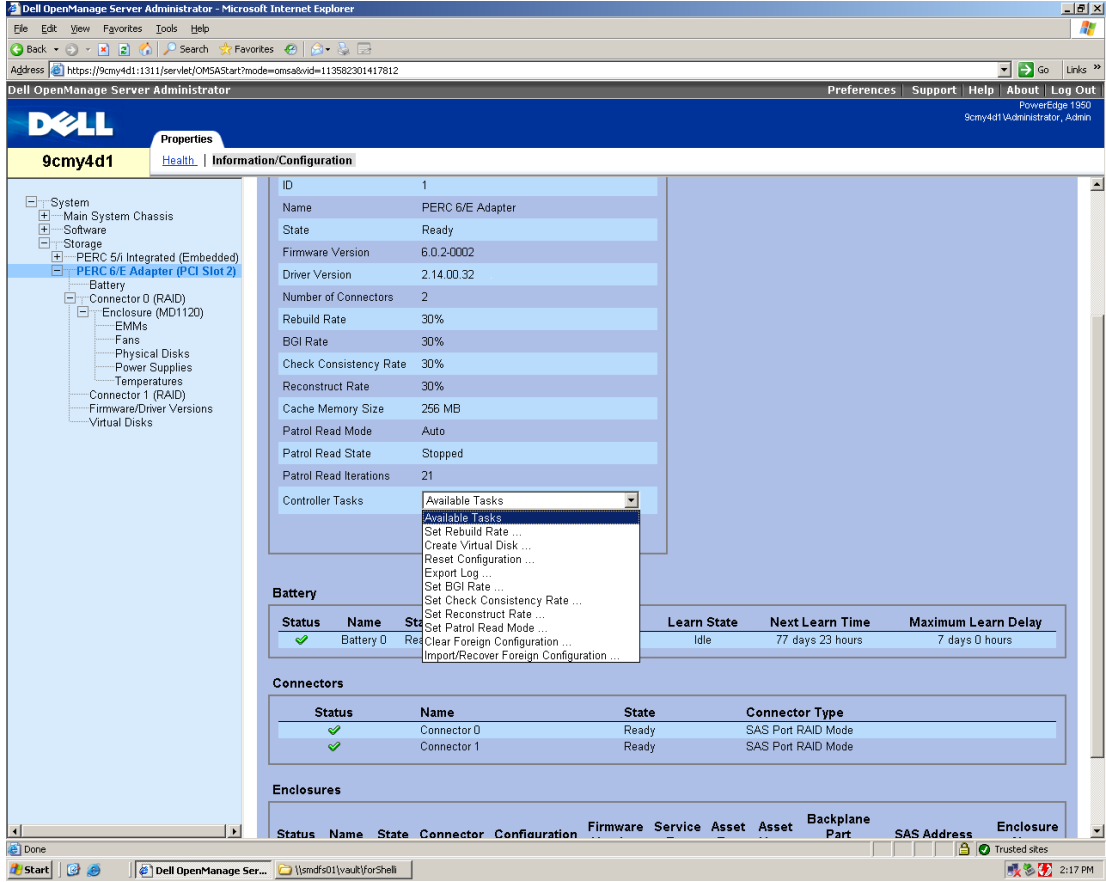
If a physical disk contains all or some portion of a foreign configuration, Server Administrator displays the physical disk state as "Foreign."

Figure 7: Physical Disks with a Foreign State



Server Administrator displays the *Import/Recover Foreign Configuration* or the *Import Foreign Configuration* task when the controller detects a foreign configuration. If you have an incomplete foreign configuration which cannot be imported, using the *Clear Foreign Configuration* task, erases the foreign data on the physical disks.

Figure 8: Controller Tasks Drop-down Menu With the Import/Recover Foreign Configuration and Clear Foreign Configuration Tasks Displayed



7. Storage Applications and Components

Identify customer usage model and needs

In order to select the correct storage solution, it is important to understand the application and user requirements. A good starting point is to use the following basic storage profiling considerations.

Table 5: Storage Profiling Considerations

Characteristic	Values	Description
Performance	<ul style="list-style-type: none"> • Throughput (MB per sec.) • I/O per Second (I/Ops) • I/O size (KB/MB) • I/O Profile (read/write and random/sequential access mix) • Latency 	Performance is the overall ability of the solution to read and write data to the disk. The performance requirements are usually determined by the type of application that is used. Different applications have different performance requirements. For example, a database or e-mail server has mostly random disk access operations while a streaming media server would have mostly sequential disk access.
Storage Capacity Needs	Gigabytes (GB) Terabytes (TB)	Storage capacity is the amount of storage space required by the application and user data. For example, e-mail storage for 100 users would require much less storage capacity than an e-mail store for 1000 users.
Storage Growth Rate	Percent increase per year	Storage growth is the expected increase in the amount of the capacity that will be required as the usage increases. This can be estimated by forecasting the number of users or clients expected to access the application in the future.
Criticality	Low, Medium, High, Very High	Criticality defines the impact to business needs if the storage is offline. This characteristic is important for choosing the right RAID level.

Picking a solution based on the application profile

One of the most important factors to consider when selecting a storage solution is the type of application that will be utilizing the storage solution. This defines the overall purpose of the server/storage solution and will determine what RAID configuration will be optimal for the application. Table 6 outlines the recommended RAID solution based on the usage of the server.

General RAID Configuration Guidelines

RAID 0

RAID 0 is generally not a recommended solution due to lack of data redundancy and protection. However, it may be utilized in situations where this is not required and maximum storage capacity and performance are essential.

RAID 1

A RAID 1 solution, while not ideal, could be utilized for small workgroup servers or servers with low storage capacity and growth requirements. Any server beyond a small workgroup server is likely to need more storage and consequently, a more efficient RAID solution.

RAID 10

RAID 10 has good I/O performance, excellent availability and redundancy. The biggest drawback of RAID 10 is that it requires twice the number of physical disks needed for data. This RAID configuration should be used in situations requiring maximum availability, redundancy, and performance.

RAID 5

RAID 5 has good I/O performance, data protection and requires only one additional physical disk than the number needed for data. RAID 5 should be used in situations where maximum storage capacity is required along with a moderate amount of data protection.

RAID 50

A balance between RAID 5 and RAID 10, this solution offers good I/O performance, availability and good storage capacity. This configuration offers slightly higher performance than RAID 5, but requires multiple additional physical disks than the number needed for data. It does however provide greater storage capacity than a RAID 10 solution with a decrease in performance and data protection. This should be used in situations where greater redundancy and data protection is required as well as a reasonable storage capacity.

RAID 6

RAID 6 has adequate I/O performance, enhanced data protection and requires only two additional physical disks than the number needed for data. RAID 6 adds an additional level of redundancy over RAID 5 with the ability to sustain 2 drive failures per RAID set. RAID 6 should be used in situations where maximum storage capacity is required along with a high amount of data protection.

RAID 60

A balance between RAID 6 and RAID 10, this solution offers good I/O performance, availability and good storage capacity. This configuration offers slightly higher performance than RAID 6, but requires multiple additional physical disks than the number needed for data. It does however provide greater storage capacity than a RAID 10 solution with a decrease in performance and data protection. This should be used in situations where greater redundancy and data protection is required as well as a reasonable storage capacity.

Note: For more details on RAID configurations see the Dell *'Getting Started with RAID'* document (<http://support.dell.com/support/edocs/storage/RAID/RAIDbk0.pdf>)

Table 6: RAID Configuration Recommendations

Application	RAID Level						
	0	1	10	5	50	6	60
E-mail	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Web	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Database or Online Transaction Processing	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Streaming Media	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
File- Archival	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
File-User File Stores	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended Not Recommended Possible

General Storage Enclosure Guidelines

MD1000 – 3U/15 3.5” Hard Drive Enclosure

The MD1000 is a storage expansion enclosure housing up to 15 3.5-inch hard drives in 3U of rack space. With its use of higher capacity 3.5” hard drives and its ability to support both high performance SAS drives as well as low cost/high capacity SATA drives, the MD1000 is the general purpose, capacity focused expansion enclosure in the Dell portfolio.

MD1120 – 2U/24 2.5” Hard Drive Enclosure

The MD1120 is a storage expansion enclosure housing up to 24 2.5-inch hard drives in 2U of rack space. With its use of 2.5” hard drives that consume less space and use less power than 3.5” hard drives, the MD1120 is the transactional performance focused expansion enclosure in the Dell portfolio.

Table 7: Storage Enclosure Recommendations

Application	Enclosure	
	MD1000 3U / 15 x 3.5” drive	MD1120 2U / 24 x 2.5” drive
E-mail	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Web	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Database or Online transaction processing	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Streaming Media	<input checked="" type="checkbox"/>	<input type="checkbox"/>
File- Archival	<input checked="" type="checkbox"/>	<input type="checkbox"/>
File-User File Stores	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended Not Recommended Possible

Application-Specific Guidelines

E-Mail Servers

The storage requirements for e-mail servers can vary depending on the size, the amount, and the type of users. While small departmental e-mail servers may work well with a small amount of storage and limited features; large corporate e-mail servers normally require greater storage capacity, very high availability, performance, and scalability. I/O profiles will vary depending on the number of users and type of mail and attachments sent.

Table 8: Email Server General I/O profile

<i>I/O Profile (Read/Write)</i>	<i>I/O Profile (Sequential/Random)</i>	<i>Bandwidth</i>	<i>I/O Size</i>	<i>Latency Sensitivity</i>	<i>Growth Rate</i>	<i>Criticality</i>
60/40	Random	Heavy	4k	High	High	High

Table 9: Email Server RAID Guidelines

<i>Application</i>	<i>RAID Level</i>							
		<i>0</i>	<i>1</i>	<i>10</i>	<i>5</i>	<i>50</i>	<i>6</i>	<i>60</i>
<i>E-Mail</i>		<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended Not Recommended Possible

Recommended:

- RAID 10 – Recommended for mission-critical e-mail servers where availability and redundancy and/or performance are the highest priority and where reduced storage capacity is acceptable.
- RAID 5 – Recommended for e-mail servers that require maximum storage capacity and moderate data protection and performance.
- RAID 50 – Recommended for those solutions that require greater redundancy and data protection and a balance between storage capacity and performance.
- RAID 6 – Recommended for e-mail servers that require maximum storage capacity and enhanced data protection and moderate performance.
- RAID 60 – Recommended for those solutions that require greater redundancy and data protection and a balance between storage capacity and performance.

Possible:

- RAID 1 – Possible solution for small e-mail servers which do not require high storage capacity.

Not recommended:

- RAID 0, Concatenated – Not recommended due to lack of redundancy and data protection.

Note: While these configurations are not recommended, they can be configured and utilized.

Table 10: Email Sever Enclosure Guidelines

Application	Enclosure	
	MD1000 3U / 15 x 3.5" drive	MD1120 2U / 24 x 2.5" drive
E-mail	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- MD1000 – Recommended for e-mail servers where large mailbox capacity is priority.
- MD1120 – Recommended for e-mail servers with typical mailbox capacities where performance and power efficiency are high priority.

Web Servers

Web servers are usually high traffic systems where read operations are the most common disk activity as web pages are requested by users. They can be intranet sites with minimal internal company traffic or global internet portal sites that receive hundreds of thousands of users daily. Performance is generally the main concern with redundancy and data protection being less of a concern. For large web sites, there are usually several web servers responding to client requests and as a result the need for redundancy and data protection is not as important. Also web servers may not require as much scalability or disk capacity since website content is relatively static and changes are usually minor content modifications.

Table 11: Web General I/O profile

I/O Profile (Read/Write)	I/O Profile (Sequential/Random)	Bandwidth	I/O Size	Latency Sensitivity	Growth Rate	Criticality
95/5	Random	Moderate	< 64k	Moderate	Low	High

Table 12: Web Server RAID Guidelines

Application	RAID Level							
		0	1	10	5	50	6	60
Database		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- RAID 10 – Recommended for enterprise web server solutions where availability and redundancy and performance are the highest priority, usually for stand-alone portal sites that are critical to the organization’s business.
- RAID 5 – Recommended for web servers that require maximum storage capacity and moderate data protection and performance.
- RAID 6 – Recommended for web servers that require maximum storage capacity and enhanced data protection and performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.
- RAID 60 – Recommended for those solutions that require a balance between storage capacity and performance with enhanced data protection.
- RAID 0 – Recommended for solutions where the web server will be part of a group of web servers that service a large internet portal and performance is the highest concern. In this situation, availability and redundancy are addressed by the cluster group and do not need to occur at the disk level.
- RAID 1 – Good solution for small websites which do not require high storage capacity.

Table 13: Web Sever Enclosure Guidelines

Application	Enclosure	
	MD1000 3U / 15 x 3.5" drive	MD1120 2U / 24 x 2.5" drive
Web Server	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- MD1120 – Recommended for web servers with typical mailbox capacities where performance and power efficiency are high priority.
- MD1000 – Recommended for e-mail servers where large mailbox capacity is priority.

Database or Online Transaction Processing Servers

Database servers can range from simple workgroup databases like Microsoft® Access™ with a few hundred users to critical Enterprise databases like Oracle or SQL Server with thousands of users. Database applications will always benefit from some data protection while other requirements, such as performance and availability, will vary. As a general rule, the more critical the database, the greater the need for data protection. Additionally, the performance requirements increase relative to the number of users accessing the database.

Online Transaction Processing (OLTP) oriented servers are used in a number of industries for the entry and retrieval of transactions. For example, OLTP is common in banking, airlines, mail-order, and supermarkets. These servers are generally critical and require maximum availability and redundancy.

Table 14: Database or OLTP General IO profile

I/O Profile (Read/Write)	I/O Profile (Sequential/Random)	Bandwidth	I/O Size	Latency Sensitivity	Growth Rate	Criticality
80/20	Random	Moderate	8k	Moderate	Low	High

Table 15: OLTP RAID Guidelines

Application	RAID Level							
	0	1	10	5	50	6	60	
OLTP	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	<input checked="" type="checkbox"/> Recommended	<input checked="" type="checkbox"/> Not Recommended		<input type="checkbox"/> Possible				

Recommended:

- RAID 10 – Since OLTP systems are critical for most business, this RAID configuration is highly recommended because of the high availability and redundancy and good performance.
- RAID 5 – Recommended for OLTP servers that require maximum storage capacity and moderate data protection and performance.
- RAID 6 – Recommended for OLTP servers that require maximum storage capacity, enhanced data protection, and moderate performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.
- RAID 60 – Recommended for those solutions that require a balance between storage capacity and performance with maximum data protection.

Possible:

- RAID 1 – Possible solution for situations which do not require high storage capacity.

Not recommended:

- RAID 0, Concatenated - These are not recommended because of the lack of redundancy and data protection.

Note: While these configurations are not recommended, they can be configured and used.

Table 16: OLTP Enclosure Guidelines

Application	Enclosure	
	MD1000	MD1120

	3U / 15 x 3.5" drive	2U / 24 x 2.5" drive
OLTP	<input type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible

Recommended:

- MD1120 – Recommended for OLTP or Database servers with typical capacities where performance and power efficiency are high priority.

Possible:

- MD1000 – Potential option for OLTP or Database servers where capacity is priority.

File Servers

File servers can be used for file archival or more dynamic storage where files are changed, added and deleted on a daily basis. These servers range from workgroup to the corporate level. Storage capacity is a key attribute of these servers as users add more and more files. File servers are generally not mission critical systems so lower levels of availability and redundancy are acceptable as data is usually backed up and can be restored in a matter of hours.

Archival file server characteristics and recommendations:

Table 17: Archival File Server General I/O profile

I/O Profile (Read/Write)	I/O Profile (Sequential/Random)	Bandwidth	I/O Size	Latency Sensitivity	Growth Rate	Criticality
90/10	Sequential	Moderate	>64k	High	Varies	Low

Table 18: Archival File Server RAID Guidelines

Application	RAID Level							
	Concatenated	0	1	10	5	50	6	60
File -Archival	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible						

Recommended:

- RAID 10 – Recommended due to high availability, redundancy, and good performance.
- RAID 5 – Recommended for file servers that require maximum storage capacity and moderate data protection and performance.
- RAID 6 – Recommended for file servers that require maximum storage capacity, enhanced data protection, and moderate performance.

- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.
- RAID 60 – Recommended for those solutions that require a balance between storage capacity and performance with maximum data protection.

Possible:

- RAID 1 – Possible solution in situations which do not require high storage capacity.

Not recommended:

- RAID 0, Concatenated - These are not recommended due to lack of redundancy and data protection. Even for long term archival storage which is intended to be a backup of important files, it is not a recommended solution.

Note: While these configurations are not recommended, they can be configured and utilized.

Table 19: Archival File Server Enclosure Guidelines

Application	Enclosure	
	MD1000 3U / 15 x 3.5" drive	MD1120 2U / 24 x 2.5" drive
File-Archival	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- MD1000 – Recommended for file servers focused on data archival due to focus on storage capacity.

Possible:

- MD1120 – Possible option for File servers where performance is high priority.

User file store characteristics and recommendations

Table 20: User Store File Server General I/O profile

I/O Profile (Read/Write)	I/O Profile (Sequential/Random)	Bandwidth	I/O Size	Latency Sensitivity	Growth Rate	Criticality
80/20	Sequential	Heavy	>64k	High	Varies	Moderate

Table 21: User Store File Server RAID Guidelines

Application	RAID Level							
	Concatenated	0	1	10	5	50	6	60
File – User file store	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- RAID 5 – Recommended for file servers that require maximum storage capacity and moderate data protection and performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.
- RAID 6 – Recommended for file servers that require maximum storage capacity, enhanced data protection, and moderate performance.
- RAID 60 – Recommended for those solutions that require a balance between storage capacity and performance with maximum data protection.

Possible:

- RAID 1 – Possible solution for situations that do not require high storage capacity.
- RAID 10 – Possible solution for situations that require greater availability and redundancy but not as much disk capacity.
- RAID 0 – While this solution is not redundant, RAID 0 may be useful for file servers that require better performance and maximum storage capacity; for example, where backups are expected to maintain data in the event of a failure.
- Concatenated – This solution could be used for situations that require maximum scalability as physical disks could be added as storage needs grow. For file servers that can afford some downtime and rely on backups to restore data, this could be a solution.

Table 22: User Store File Server Enclosure Guidelines

Application	Enclosure	
	MD1000 3U / 15 x 3.5" drive	MD1120 2U / 24 x 2.5" drive
File-User File Store	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- MD1000 – Recommended for user file store file servers in which capacity is highest priority.
- MD1120 – Recommended for user file store file servers where performance and power efficiency are high priority.

Streaming Media Servers

Streaming media servers are systems that provide web-casting, video conferencing, Internet entertainment (for example, Internet TV or Internet radio), and multimedia services. These systems generally require a balance between storage capacity, availability, redundancy, and performance. Like web servers, they may also be part of a group of systems that work together to provide content.

Table 23: Streaming Media IO profile

<i>I/O Profile (Read/Write)</i>	<i>I/O Profile (Sequential/Random)</i>	<i>Bandwidth</i>	<i>I/O Size</i>	<i>Latency Sensitivity</i>	<i>Growth Rate</i>	<i>Criticality</i>
98/2	Sequential	Heavy	>64k	High	High	Varies

Table 24: Streaming Media RAID Guidelines

<i>Application</i>	<i>RAID Level</i>							
	<i>Concatenated</i>	<i>0</i>	<i>1</i>	<i>10</i>	<i>5</i>	<i>50</i>	<i>6</i>	<i>60</i>
Streaming Media	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	<input checked="" type="checkbox"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible					

Recommended:

- RAID 5 – Recommended for streaming media servers that require maximum storage capacity and moderate protection and performance. This is the recommended RAID configuration for stand alone streaming media servers.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance. This solution provides greater availability and redundancy but less storage capacity.
- RAID 6 – Recommended for streaming media servers that require maximum storage capacity, enhanced protection, and moderate performance.
- RAID 60 – Recommended for those solutions that require a balance between storage capacity and performance with maximum data protection. This solution provides greater availability and redundancy but less storage capacity.
- RAID 0 – While not redundant, this solution may be useful for streaming media servers that require better performance and are members of a cluster of servers where availability and redundancy are handled at a higher level.

Possible:

- RAID 1, RAID 10 – Possible in situations which do not require high storage capacity. In general servers utilized for streaming media require large amounts of storage capacity for content, especially uncompressed video.

Not Recommended:

- Concatenated – This solution is not recommended due to lack of redundancy and limited performance.

Note: While these configurations are not recommended, they can be configured and utilized.

Table 25: Streaming Media Enclosure Guidelines

Application	Enclosure	
	MD1000 3U / 15 x 3.5" drive	MD1120 2U / 24 x 2.5" drive
Streaming Media	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Recommended
 Not Recommended
 Possible

Recommended:

- MD1000 – Recommended for streaming media servers due to focus on storage capacity.

Possible:

- MD1120 – Possible option for streaming media servers where transactional performance is high priority.

i.

Hot Spares

Hot Spare functionality provides extra security, availability and redundancy by automatically replacing a failed physical disk in a RAID group and allowing the rebuild of the degraded array to begin immediately. PERC controllers support Hot Spare functionality and it is recommended whenever possible.

8. Summary and Conclusions

It is important to understand the characteristics of each storage component in order to assess the performance of the overall application. It is also necessary to understand how the storage component is configured since this has a direct impact on the effectiveness of an application. Table 26 summarizes typical storage recommendations for various applications. Recommendations listed in the table are general conclusions and the actual storage solution may have additional factors which need to be considered before determining an optimal solution. Additional factors could include:

- Data and power back-up strategies of an Enterprise IT
- Storage deployment environment:
 - Number of active and passive users
 - Type of data – mission critical or non-mission critical
- Storage deployment budget
- Long term storage consolidation strategy

Table 26: Typical Storage Recommendation for Different Storage Applications

Application		Storage Recommendations			Remarks
Servers	Key Parameters	RAID Level	PowerVault MD1000	Host Controller	
Email	Availability Performance Scalability	RAID 10 provides both availability and performance	Daisy-chaining provides capacity scalability	For added performance enable Write back cache	Recommendation targets a typical enterprise e-mail Server
Database	Availability Performance	RAID 10 provides both availability and performance	Daisy-chaining provides capacity scalability	For added performance enable Write back cache	Performance requirements for database servers depend on number of users
Web	Performance	RAID 0 in Web Server farms provides maximum performance of all RAID levels	Split Topology spanning RAID 0 across channels provides better performance if daisy-chaining for additional capacity is not required	Enable read ahead capabilities for added performance boost	Web server farms, usually have multiple servers providing replicated web data. Thus, data availability is handled at server level
On-line Transaction	High Availability	RAID 10 provides highest availability of all RAID levels	Daisy-chaining provides capacity scalability	For added performance enable Write back cache	Recommendation targets Enterprise applications which require higher availability than performance. In general for all OLTP, data protection is of utmost importance
File	Scalability	RAID 5 maximizing storage with some data protection	Daisy-chaining provides capacity scalability	Enable Global Hot Spare for improved availability for RAID 5	These servers usually have data backed up periodically and hence, require minimum data protection from storage components
Streaming Media	Performance	RAID 0 provides maximum performance of all RAID levels	Split Topology spanning RAID 0 across channels provides better performance if daisy-chaining for additional capacity is not required	Enable read ahead capabilities for added performance boost	Performance is of utmost importance for Streaming servers. Data on these servers is usually backed up on external media such as tape libraries or optical media - CD/DVD

Appendix – A: References

1. SAS Protocol [Specification](#)
(http://www.t10.org/drafts.htm#SCSI3_SAS)
2. [SFF-8470 Connector Specification](#)
(<http://www.sffcommittee.com/ns/index.html>)
3. [Beyond The 2-TB SCSI Logical Unit](#)
(http://www.dell.com/downloads/global/vectors/2004_2tblun.pdf)
4. [Dell PowerVault MD1000 Documentation](#)
(<http://support.dell.com/support/edocs/systems/md1000/>)
5. [Dell PowerVault MD1120 Documentation](#)
(<http://support.dell.com/support/edocs/systems/md1120/>)
6. [Dell Enterprise Technology Center - Storage](#)
<http://www.delltechcenter.com/page/Storage>

9. Appendix – B: Glossary

Array – Collection of multiple physical disks to form a single logical volume

BBU – Battery Backup Unit

BGI – BackGround Initialization

BIOS – Basic Input Output System

Cluster – Multi-initiator Environment with MSCS operating environment

HBA – Host Bus Adapter (non-RAID controllers)

Hot Plug – Insertion or removal of device without the need to quiescent I/Os

JBOD – Industry term used to describe a storage enclosure without integrated RAID capabilities.
Derived from the expression “Just a Bunch Of Disk”

MD1xxx Family – The Dell PowerVault MD1xxx Family include the MD1000 and MD1120 enclosures.

MD3xxx Family – The Dell PowerVault MD3xxx Family include the MD3000 and MD3000i storage arrays

PERC – PowerEdge Expandable RAID Controller

RAID – Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks

SAS – Serial Attached SCSI

SATA – Serial Advanced Technology Attachment

SCSI – Small Computer System Interface

SES – Storage Enclosure Services

U320 – Ultra 320 (320MB/sec through put)

10. Appendix – C: Revision History

<i>Revision</i>	<i>Description</i>	<i>Date</i>	<i>Modified By</i>
A00	First Release	4/17/2006	SAS Solution Core Team
A02	Updated to reflect Dell PowerVault MD1000 and MD1120 enclosures (aka. MD1xxx Family).	6/16/2008	SAS Solution Core Team