



A Reference Guide for Optimizing Dell™ SCSI Solutions

Authored By:
Dell SCSI Storage Solution Team

November 17, 2005 rev A02

Information in this document is subject to change without notice.

© Copyright 2005 Dell Inc. All rights reserved.

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

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Dell, the *Dell* Logo, *PowerVault*, *PowerEdge* and *OpenManage* are trademarks of Dell Inc. *Microsoft* is a registered trademark of Microsoft Corporation. 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

A REFERENCE GUIDE FOR OPTIMIZING DELL SCSI SOLUTIONS	3
1. INTRODUCTION	3
2. STORAGE APPLICATIONS AND COMPONENTS	4
<i>Identify customer usage model and needs</i>	4
<i>Picking a solution based on the application profile</i>	4
<i>General RAID Configuration Guidelines</i>	5
<i>Application specific guidelines</i>	6
<i>E-Mail Servers</i>	6
<i>Database Servers</i>	7
<i>Web Servers</i>	8
<i>Online Transaction Processing Servers</i>	9
<i>File Servers</i>	10
<i>Streaming Media Servers</i>	12
<i>Additional configuration factors</i>	13
<i>Hot Spares</i>	13
<i>Clustering Considerations</i>	13
<i>Considerations for Optimal RAID Performance</i>	13
3. DELL'S PRODUCT OFFERING	14
<i>PowerVault™ 220S & 221S Storage Enclosure</i>	14
<i>PERC – PowerEdge™ RAID Controllers</i>	14
<i>SCSI Cabling</i>	15
<i>SCSI Hard Drive Support</i>	16
<i>3rd Party Hardware support</i>	16
<i>Server Support</i>	16
4. PV22xS ENCLOSURE OVERVIEW	19
<i>Joined Bus Topology</i>	19
<i>Split Bus Topology</i>	20
<i>Cluster Topology</i>	21
<i>PV22xS – Common Features and Capabilities</i>	22
<i>SCSI ID Assignment</i>	22
<i>Multiple PV22xS Enclosures on a Single Host System</i>	23
<i>Mixing EMM Speeds</i>	23
<i>Mixing HDD Capacities</i>	24
<i>Mixing HDD Spindle Speeds</i>	24
<i>SCSI Enclosure Services (SES)</i>	24
<i>Warm and Hot Pluggable Features</i>	24
<i>External SCSI Cables</i>	25
<i>External SCSI Connector Types</i>	25
<i>PV22xS Configuration Topology and Application Attributes</i>	26
5. SCSI HOST BUS ADAPTERS / SCSI RAID	29
<i>RAID: Hardware Vs Software</i>	29
<i>Considerations to Implement a RAID Solution</i>	29
<i>Controller Family</i>	30
<i>Data Protection Strategy</i>	32
<i>High Availability Needs – Clustering</i>	33
<i>Heterogeneous vs. Homogeneous environments</i>	33
<i>Interface Channel Options – Internal/External/Multi channel device configuration</i>	34
<i>Capacity Needs – Virtual Disk Size</i>	35
<i>Additional RAID information</i>	35
6. SUMMARY AND CONCLUSIONS	39
7. APPENDIX – A: REFERENCES	41
8. APPENDIX – B: GLOSSARY	42
9. APPENDIX – C: ABOUT THE DELL SCSI STORAGE SOLUTION TEAM	43
<i>Authors</i>	43
10. APPENDIX – D: REVISION HISTORY	44

A Reference Guide for Optimizing Dell SCSI Solutions

1. Introduction

Enterprise storage applications involve interconnection of multiple components – including both software and hardware aspects which would complete a storage solution. Any given storage application may be supported by one or more storage solutions. Thus, it becomes essential to identify an optimum storage solution, one which would be ideal for a given application.

Selection of storage components not only impacts the effectiveness of the solution for an application but also would impact the performance of Enterprise IT infrastructure which relies on that storage. Key storage components which impact a storage solution can be categorized as:

- Physical Storage Device Technology – PowerVault™ Direct Attached Storage PV22xS family
- Interconnect Technology – Ultra160/Ultra320 SCSI
- Host System Interface Technology – non-RAID (HBA) and RAID controllers (PERC)
- Management Capabilities – OpenManage™ Storage Management (OMSM) or Array Manager

While selection of storage components is essential for building a storage solution, on the other hand it is also critical to understand the requirements for a given application. In general, effectiveness of all the applications is driven by the following key storage attributes and the usage model for the application, which should then determine the relative importance for each of these attributes:

- *Reliability* – Storage solution provides reliable access to data
- *Availability* – Data from storage devices readily available
- *Serviceability* – In case of failure of storage components, solution should allow easy access to assemblies/components and be easy to service
- *Redundancy* – Solution provides redundancy for both the data storage, as well as for hardware management capabilities
- *Performance*
 - Bandwidth Utilization – Data Throughput/Available Bandwidth
 - I/Os per second (IOPS)
- *Data Protection* – Data is protected against catastrophic failures on a storage or system component
- *Scalability* – Storage capacity can be expanded as applications necessitate
- *Cost* – driven by RAID and other hardware configurations.

An optimum storage solution would be something comprised of storage components that would address all the key attributes for a given application, keeping their relative importance in mind. This paper presents a study of all storage components from technological to functional, as well as behavioral differences, to help define an optimum storage solution. In this paper, all the storage components are measured against the attributes listed above which impact the overall effectiveness of an application.

Understanding each of the storage components with respect to these attributes will help in making intelligent decisions to determine an optimal configuration for a given application.

The paper presumes that the reader has a basic understanding of different levels of RAID and interconnects technologies – Parallel SCSI (SPI-4), PCI vs. PCI-X vs. PCI-e. This paper presents a study of all storage components, as identified earlier, against various storage applications and different operating environments; however, specific behavior of each of the storage applications and management functions is beyond the scope of this document. In this document we present storage solution from Parallel SCSI perspective, specifically for Dell’s Direct Attached Storage (DAS) products. Other storage technologies such as SAS, SATA, or Fibre Channel, as well as topologies such as Network Attached Storage (NAS) or Storage Area Network (SAN) are outside the scope of this document. All storage applications referred to herein are treated independent of any specific host system or server.

2. Storage Applications and Components

Identify customer usage model and needs

In order to select the correct storage solution for any given situation, it is important to understand what the application and user requirements will be for that solution. A good starting point is to have an understanding of basic storage profiling considerations.

• Table 2-1: Storage Profiling Considerations

<i>Characteristic</i>	<i>Values</i>	<i>Description</i>
<i>Performance</i>	<ul style="list-style-type: none"> • Bandwidth (MB per sec.) • 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 disk. The performance requirements are usually determined by the type of application being utilized. 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.
<i>Storage Capacity Needs</i>	Gigabytes	Storage capacity is the current 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.
<i>Storage Growth Rate</i>	Percent increase per year	Storage growth is the expected increase in the amount of capacity that will be required as usage of the server increases. This is usually measured in the expected number of users or clients accessing the server
<i>Criticality</i>	Low, Medium, High, Very High	Criticality defines the impact to mission critical business needs if the storage is offline. This characteristic is important for choosing the right RAID level, and determining if clustering is needed.

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 the most optimal for the application. Table 2-2 outlines the recommended RAID solution based on the usage of the server.

• Table 2-2: Application Vs RAID configurations

Application	RAID Level					
	Concatenated	0	1	10	5	50
Email	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Database	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Web	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
On-line transaction	<input checked="" type="checkbox"/>					
File - Archival	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
File - User File stores	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input 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"/>

Recommended Not Recommended Possible

General RAID Configuration Guidelines

RAID 0

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

RAID 1

A RAID 1 solution, while not ideal for most servers, could be utilized for small workgroup servers or servers with low storage capacity and growth requirements. Any servers beyond a small workgroup server are 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 is the higher relative cost for significantly lower storage capacity. This RAID configuration should be used in situations where the maximum availability, redundancy and performance are the priority considerations.

RAID 5

RAID 5 has good I/O performance, data protection and relative cost with excellent storage capacity. 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 at a slightly higher cost and reduced storage capacity. 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 reasonable storage capacity.

Note: For more details on RAID configurations refer to the Dell™ ‘Getting Started with RAID’ document (<http://support.dell.com/support/edocs/storage/RAID/RAIDbk0.pdf>).

Concatenated Container

This solution is not recommended due to lack of data protection and redundancy and no performance gain. However for non-critical implementations that require high scalability this type of solution will provide more ease of use than other RAID configurations.

Application specific guidelines

E-Mail Servers

The storage requirements for e-mail servers can vary depending on the size, the amount, and 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. IO profiles will vary depending on the number of users and type of mail and attachments sent.

• Table 2-3: 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 2-4: Email Sever 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>E-Mail</i>	<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"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible			

Recommended:

- RAID 10 – Recommended for enterprise mission critical e-mail servers where availability and redundancy and/or performance are the highest priority and reduce storage capacity is acceptable.
- RAID 5 – Recommended for e-mail servers that require maximum storage capacity and only base 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.

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.

Database Servers

Database servers can range from simple workgroup databases like Microsoft® Access with a few hundred users to mission 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 mission critical the database is, the more data protection is required. Additionally, the performance requirements increase relative to the number of users accessing the database.

• Table 2-5: Database General IO profile

<i>I/O Profile (Read/Write)</i>	<i>I/O Profile (Sequential/Random)</i>	<i>Bandwidth</i>	<i>IO Size</i>	<i>Latency Sensitivity</i>	<i>Growth Rate</i>	<i>Criticality</i>
70/30	Random	Heavy	8K	High	High	High

• Table 2-6: Database sever 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>Database</i>	<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"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible			

Recommended:

- RAID 10 – Recommended for enterprise mission critical database solutions where availability and redundancy and/or performance are the highest priority.

- RAID 5 – Recommended for databases that require maximum storage capacity and only base data protection and performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.

Possible:

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

Not recommended:

- RAID0, Concatenated - Not recommended due to lack of redundancy and data protection

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

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 moderate intranet sites with minimal traffic 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 a concern. For large web sites, usually there are 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 2-7: Web General IO profile

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

• Table 2-8: Web 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>Database</i>	<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 – Recommended for enterprise web server solutions where availability and redundancy and performance are the highest priority, usually stand-alone portal sites that is critical to the organization’s business.

- RAID 5 – Recommended for web servers that require maximum storage capacity and only base data protection and performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.
- RAID 0 – Recommended for solutions where 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 is handled by the cluster group and does not need to happen at the disk level.
- RAID 1 – Good solution for small websites which do not require high storage capacity.

Not recommended:

- Concatenated - This solution is not recommended due to lack of redundancy and data protection

Note: While this configuration is not recommended, it can be configured and utilized.

Online Transaction Processing Servers

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. They are generally mission critical servers that require the maximum availability and redundancy possible.

• Table 2-9: OLTP General IO profile

<i>I/O Profile (Read/Write)</i>	<i>I/O Profile (Sequential/Random)</i>	<i>Bandwidth</i>	<i>IO Size</i>	<i>Latency Sensitivity</i>	<i>Growth Rate</i>	<i>Criticality</i>
80/20	Random	Moderate	2k-8k	Moderate	Low	High

• Table 2-10: OLTP 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>OLTP</i>	<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"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible			

Recommended:

- RAID 10 – Since OLTP systems are very critical for most business, this RAID configuration is highly recommended due to the high availability and redundancy and good performance.
- RAID 5 – Recommended for OLTP servers that require maximum storage capacity and only base data protection and performance.

- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.

Possible:

- RAID 1 – Possible solution for 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.

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

File Servers

File servers can be archival long term storage repositories or more dynamic user file storage where files are changed, added and deleted on a daily basis. They can range from workgroup to company level in their scope. A key factor for file servers is storage capacity as users add more files. File servers are generally not mission critical systems so reduced availability and redundancy or none at all is acceptable as the data is usually backed up and can be restored in a matter of hours.

Archival file server characteristics and recommendations

• Table 2-11 Archival File Server General I/O profile

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

• Table 2-12: Archival File Server 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>File - Archival</i>	<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"/> Recommended	<input checked="" type="checkbox"/> Not Recommended	<input type="checkbox"/> Possible			

Recommended:

- RAID 10 – Recommended due to the high availability and redundancy and good performance.
- RAID 5 – Recommended for file servers that require maximum storage capacity and only base data protection and performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.

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. Since long term archival storage is indented in some ways 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.

User file store characteristics and recommendations

• Table 2-13 User Store File Server General IO profile

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

• Table 2-14: User Store File Server 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>File – User file stores</i>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input 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 file servers that require maximum storage capacity and only base data protection and performance.
- RAID 50 – Recommended for those solutions that require a balance between storage capacity and performance.

Possible:

- RAID 1 – Possible solution situations which 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 fault tolerant, 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 use for situations that require maximum scalability as drives 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.

Streaming Media Servers

Streaming media servers are systems that provide web casting, video conferencing, internet entertainment like TV or radio, and multimedia content oriented services. These are systems that 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 2-15 Database General IO profile

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

• Table 2-16: Database sever 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>Streaming Media</i>	<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"/> 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 only base data protection and performance. This would be 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 would provide greater availability and redundancy but less storage capacity.
- RAID 0 – While this solution is not redundant, this may be useful for streaming media servers that require better performance and are members of a cluster of servers where availability and redundancy is handled at a higher level.

Possible:

- RAID 1, RAID 10 – Possible solutions 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.

Additional configuration factors

Hot Spares

Hot Spare functionality provides extra security availability and redundancy by replacing a failed drive, allowing the rebuild of the degraded array to begin immediately. Both the Adaptec and LSI based PERC controllers' support this feature, but a simple SCSI controller would not. See the controller user's guide for details and limitations. Use of Hot Spares is recommended whenever possible.

Clustering Considerations

In a cluster environment the cache being disabled on the controller poses a significant performance impact to RAID 5 and RAID 50 configurations. If the system is to be clustered RAID 10 is the optimal RAID level.

• Table 2-17: Cluster Solution Vs RAID Levels

Cluster Solutions	RAID Level					
	Concatenated	0	1	10	5	50
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	<input checked="" type="checkbox"/> Not Recommended		<input checked="" type="checkbox"/> Recommended		<input type="checkbox"/> Possible	

Considerations for Optimal RAID Performance

The PowerVault™ 22xS is capable of offering increased performance by allowing multiple channels to connect to the storage enclosure. This ability can increase the performance of most configurations by allowing a multi channel RAID controller to increase the bandwidth of the SCSI bus. For more information see Table 4-2: PV22xS Configuration Topologies vs RAID Level Performance.

• Table 2-12-7: RAID Level Benefit (Generalized)

	Clustering	Multi Channel	Relative Cost	Capacity	Scalability
Concatenated drives	None	None	None	Excellent	Excellent
RAID 0	None	Excellent	None	Excellent	None
RAID 1	Moderate	Excellent	Excellent	None	None
RAID 10	Excellent	Excellent	Excellent	None	Excellent
RAID 5	Moderate	Moderate	Moderate	Moderate	Moderate
RAID 50	Moderate	Moderate	Moderate	Moderate	Moderate

3. Dell's Product Offering

PowerVault™ 220S & 221S Storage Enclosure

Dell™ PowerVault 220xS systems are flexible, external SCSI expansion enclosures designed to support multiple Dell storage environments and RAID configurations. Each system offers maximized drive-spindle count, hot-plug hard drives, optional redundant power, redundant cooling, rack mount capability, systems management features, and a modular design for easy upgrades. Most major components, including hard drives and power supply/cooling modules are hot-pluggable and can be removed and replaced easily.

The PV22xS is available as two model types: the PV220S which fits within an industry standard 19 inch rack, and the PV221S which is a free standing floor tower. Both models leverage the same base mechanical frame, but differ with respect to the mounting hardware and front bezel features. Dell offers a conversion kit to change from a PV220S to a PV221S. Basic features and options common to both models are shown in Figure 3-1.

PowerVault 220S/221S Storage Arrays

Performance

- U320 SCSI speed
- dual channel PCI-Express PERC 4 U320 controller

Capacity

- 4.2TB with 300GB hard drives
- Up to 32 disks per virtual disk

Scalability

- Up to eight enclosures per server
- Scales up to 32TB per server

Manageability

- Robust RAID management
- Fully integrated solution

Affordability

- Lower cost than Fibre Channel
- Low cost non-redundant configs



PV 220S Rack SCSI Storage Array



PV 221S Tower SCSI Storage Array

FEATURES

- Up to 14 1" SCSI U160 SCSI hard disk drives
- 73GB, 146 & 300GB 10K RPM or 18GB, 36GB & 73GB 15K RPM disks
- Supports RAID 0, 1, 5, 10, 50 with PERC 3 or PERC 4 RAID cards
- 3+1 Fan Redundancy, optional redundant power supply
- Supports cables up to 12 meters from host server
- Configurable in joined bus, split bus or cluster mode
- Leverages PERC 3 U160 or PERC 4 RAID controller card
 - Allows 32 hard disk drive per virtual disk
 - Up to forty (40) virtual disks per controller
 - Dual channel and Quad Channel
- Host clustering support (Microsoft Cluster Support)
- Externally accessible Ultra3 (Ultra160) or Ultra 320 (U320) SCSI
- 3U form factor for rack dense environments

Figure 3-1 PV22xS Enclosure Features and Options

PERC – PowerEdge™ RAID Controllers

The PowerEdge RAID Controller (PERC) is a high-performance, intelligent peripheral component interconnect (PCI) and PCI-Express to Small Computer System Interface

(SCSI) host adapter with RAID control capabilities. It provides reliable fault-tolerant disk subsystem management and is an ideal RAID solution for internal storage in Dell's PowerEdge enterprise systems. The RAID controller offers a cost-effective way to implement RAID in a server.

Dell is currently offering its fourth generation of RAID controllers, the PERC 4e family, supporting the industry standard Ultra320 SCSI storage interconnect. PERC 4e RAID controllers support the PCI-Express host interconnects for maximum bandwidth connectivity & expandability. The PERC 4e family of controllers includes the PERC 4e/SI and PERC 4e/DI for RAID On Motherboard solutions and the PERC 4e/DC add-in HBA. The following matrix describes the benefits of each solution.

Previous generations of PERC have included the PERC 2, PERC 3 and PERC 4 families. The following PERC solutions support the PowerVault 220S and 221S External SCSI Storage solution:

• Table 3-1: PERC Solution supported by PV22xS

Maximum Speed	Card
80 MB/s	PERC 2/DC PERC 2/QC
160 MB/s	39160 HBA PERC 3/DC PERC 3/QC
320 MB/s	PE 2600 ROMB (PERC 4/Di on this server) PE 1750 ROMB (PERC 4/Di on this server) PERC 4/SC PE 2800 ROMB (PERC 4e/Di on this server) PE 2850 ROMB (PERC 4e/Di on this server) PE 6800 ROMB (PERC 4e/Di on this server) PE6850 ROMB (PERC 4e/Di on this server) PERC 4/DC (PCI, dual channel U320 RAID add-in HBA) PERC 4e/DC (PCI-Express, dual channel U320 RAID add-in HBA)

SCSI Cabling

HBA and PERC controllers support a low-voltage differential (LVD) SCSI bus. LVD enables cables up to 12 meters long. Throughput on each SCSI channel can be as high as 320 MB/sec.

• Table 3-2: PERCs and Cable Type

PV22xS Controller/Cable Matrix	
Controller type	Cable type
PERC Add-in Cards	
PERC 2/DC	Wide HD to Slim VHDCI
PERC 2/QC	Wide HD to Slim VHDCI
PERC 3/QC	Wide HD to Slim VHDCI
PERC 3/DCL	Wide HD to Slim VHDCI
PERC 3/Di Embedded	Wide HD to Wide HD on PE2600, PE2500, and 4400 Slim VHDCI to Wide HD on PE1650, PE1750
PERC 4/DC	Wide HD to Slim VHDCI
PERC 4/Di Embedded	Wide HD to Wide HD on PE2600
PERC 4e/DC	Wide HD to Slim VHDCI
PERC 4e/Di Embedded	Wide HD to Wide HD
SCSI HBAs	
Adaptec 3960D (39160)	Wide HD to Wide HD

SCSI Hard Drive Support

PowerVault disk-based storage solutions (such as the PowerVault 220S) support the latest SCSI hard drives, currently specified as Ultra 320 SCSI. SCSI solutions from Dell support both 10,000 RPM and 15,000 RPM spindle speeds and several variations of capacities to allow customers the maximum degree of flexibility. For the latest information of support for specific hard drives, please visit the Dell Storage web site at:

<http://www1.us.dell.com/content/products/category.aspx/storage>.

3rd Party Hardware support

Dell development teams spend a significant amount of time testing and qualifying SCSI hard drives and PERC controllers with PowerVault 22xS, and PERC controllers (as well as PowerEdge servers), to ensure the highest quality products. As a result of this testing, hard drives purchased direct from Dell can include upgrades not generally available to the rest of the industry. As such, Dell supports only Dell tested and qualified hardware. In order to help ensure the best possible customer experience, Dell highly recommends purchasing all hardware direct from Dell for the PowerVault 22xS, PowerEdge servers or any other Dell products. All hardware (disk drives, RAID controllers, cables, power supplies) purchased direct from Dell carries the Dell limited warranty¹ and helps ensure full compatibility. Materials purchased from 3rd party vendors for use in Dell systems are not protected by a Dell warranty and may not deliver an acceptable customer experience.

Server Support

Dell PowerVault SCSI Solutions support all PowerEdge, PowerVault NAS and many PowerEdge SC servers. This adds tremendous value by delivering a single integrated Enterprise solution that leverages industry standard hardware platforms and a common management framework for customers' server and storage solutions.

¹For a copy of our guarantees or limited warranties, please write Dell USA L.P., Attn: Warranties, One Dell Way, Round Rock, TX 78682. For more information, visit www.dell.com/warranty.

In addition, many Dell servers & storage solutions offer cross generational support to ensure customers have a seamless experience when adopting new technologies. The following compatibility matrix indicates the cross-generational support Dell offers for its SCSI Solutions:

• Table 3-3: PERC 4 Support Matrix (6G, 7G & 8G servers)

Y = Supported, N = Not Supported, # = maximum # of factory installed cards
 FI = Factory install; CK = Customer Kit

System	4/Di (ROMB)		4/Si (ROMB)		4/IM		4/SC		4/DC		4e/DC		4e/Di (ROMB)		4e/Si (ROMB)		
	FI	CK	FI	CK	FI	CK	FI	CK	FI	CK	FI	CK	FI	CK	FI	CK	
PowerEdge Servers																	
PE 400SC	N		N		N		N		N		N		N		N		N
PE 420SC	N		N		N		N		N		N		N		N		N
PE 300	N		N		N		N		N		N		N		N		N
PE 350	N		N		N		N		N		N		N		N		N
PE 500SC	N		N		N		N		N		N		N		N		N
PE 600SC	N		N		N		N		N		N		N		N		N
PE 650	N		N		N		1	Y	1	Y	N		N		N		N
PE 700	N		N		N		1	Y	N		N		N		N		N
PE 750	N		N		N		1	Y	1	Y	N		N		N		N
PE 800	N		N		N		N		1	Y	N		N		N		N
PE 1400	N		N		N		N		N		N		N		N		N
PE 1400SC	N		N		N		1	Y	N		N		N		N		N
PE 1420SC	N		N		N		N		N		N		N		N		N
PE 1500	N		N		N		1	Y	N		N		N		N		N
PE 1550	N		N		N		N		N		N		N		N		N
PE 1600SC	N		N		N		N		N		N		N		N		N
PE 1650	N		N		N		N		1	Y	N		N		N		N
PE 1655MC	N		N		N		N		N		N		N		N		N
PE 1855MC	N		N		N		N		N		N		N		N		N
PE 1750	Y		N		N		N		1	Y	N		N		N		N
PE 1800	N		N		N		1	Y	1	Y	1	Y	Y		Y		Y
PE 1850	N		Y		N		1	Y	1	Y	1	Y	N		N		Y
PE 2400	N		N		N		N		N		N		N		N		N
PE 2450	N		N		N		N		N		N		N		N		N
PE 2500	N		N		N		N		2	Y	N		N		N		N
PE 2550	N		N		N		N		2	Y	N		N		N		N
PE 2600	Y		N		N		N		2	Y	N		N		N		N
PE 2650	N		N		N		N		2	Y	N		N		N		N
PE 2800	1		N		N		N		2	Y	2	Y	Y		Y		N
PE 2850	1		N		N		N		2	Y	2	Y	Y		Y		N
PE 2855	N		N		Y		N		N		N		N		N		N
PE 3250	N		N		N		1	Y	2	Y	N		N		N		N
PE 7250	N		N		N		N		N		N		N		N		N
PE 4400	N		N		N		N		2	Y	N		N		N		N
PE 4600	N		N		N		N		2	Y	N		N		N		N
PE 6400	N		N		N		N		3	Y	N		N		N		N
PE 6450	N		N		N		N		3	Y	N		N		N		N
PE 6600	N		N		N		N		3	Y	N		N		N		N
PE 6650	N		N		N		N		3	Y	N		N		N		N
PE 6800	1		N		N		N		3	Y	3	Y	Y		Y		N
PE 6850	1		N		N		N		3	Y	3	Y	Y		Y		N
PE 7150	N		N		N		N		N		N		N		N		N
PE 8250	N		N		N		N		3	Y	N		N		N		N
PowerVault NAS																	
PV 725N	N		N		N		N		N		N		N		N		N
PV 735N	N		N		N		N		N		N		N		N		N
PV 745N	N		N		N		1	Y	1	Y	N		N		N		N
PV 750N	N		N		N		N		N		N		N		N		N
PV 755N	N		N		N		N		N		N		N		N		N
PV 770N	Y		N		N		N		2	Y	N		N		N		N
PV 775N	N		N		N		N		2	Y	N		N		N		N
PV220 Support																	
PV22xS	Y		Y		N		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

• Table 3-4: PERC 2 & PERC 3 Support Matrix (4G & 5G servers)

Y = Supported, N = Not Supported, # = maximum # of factory installed cards
 FI = Factory install; CK = Customer Kit

System	2/SC		2/DC		3/DI		3/SC		3/DCL		3/DC		3/QC	
							FI	CK	FI	CK	FI	CK	FI	CK
PowerEdge Server Support														
PE 500SC	N	N	N	N	N	N	1	Y	N	N	N	N	N	N
PE 600SC	N	N	N	N	N	N	1	Y	N	N	N	N	N	N
PE 1400	N	N	N	N	N	N	1	Y	1	Y	N	N	N	N
PE 1400SC	N	N	N	N	N	N	1	Y	1	Y	N	N	N	N
PE 1500	N	N	N	N	N	N	1	Y	N	N	N	N	N	N
PE 1550	N	N	N	N	N	N	N	N	1	Y	1	Y	N	N
PE 1600SC	N	N	N	N	N	N	1	Y	N	N	N	N	N	N
PE 1650	N	N	N	N	1	Y	N	N	N	N	1	Y	N	N
PE 2400	N	N	1	Y	N	N	N	N	N	N	2	Y	2	Y
PE 2450	N	N	1	Y	N	N	N	N	N	N	2	Y	2	Y
PE2500	N	N	N	N	1	Y	N	N	N	N	2	Y	2	Y
PE 2550	N	N	N	N	1	Y	N	N	N	N	2	Y	2	Y
PE 2600	N	N	N	N	N	N	N	N	N	N	2	Y	2	Y
PE 2650	N	N	N	N	1	Y	N	N	N	N	2	Y	2	Y
PE 3250	N	N	N	N	N	N	N	N	N	N	2	Y	N	N
PE 7250	N	N	N	N	N	N	N	N	N	N	N	N	N	N
PE 4400	N	N	1	Y	N	N	N	N	N	N	2	Y	2	Y
PE 4600	N	N	N	N	1	Y	N	N	N	N	2	Y	2	Y
PE 6400	N	N	1	Y	N	N	N	N	N	N	3	Y	2	Y
PE 6450	N	N	1	Y	N	N	N	N	N	N	3	Y	2	Y
PE 6600	N	N	N	N	N	N	N	N	N	N	3	Y	2	Y
PE 6650	N	N	N	N	N	N	N	N	N	N	3	Y	2	Y
PE 7150	N	N	N	N	N	N	N	N	N	N	3	Y	2	Y
PE 8250	1	Y	1	Y	N	N	N	N	N	N	3	Y	2	Y
PowerVault NAS Support														
PV 725N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
PV 735N	N	N	N	N	N	N	N	N	N	N	1	Y	2	Y
PV 750N	N	N	N	N	N	N	N	N	N	N	2	Y	2	Y
PV 755N	N	N	N	N	N	N	N	N	N	N	1	Y	2	Y
PV 770N	N	N	N	N	N	N	N	N	N	N	2	Y	2	Y
PV 775N	N	N	N	N	N	N	N	N	N	N	1	Y	2	Y
PowerVault Enclosure Support														
PV20xS	Y	Y	Y	Y	N	N	N	N	Y	Y	Y	Y	Y	Y
PV21xS	Y	Y	Y	Y	N	N	N	N	Y	Y	Y	Y	Y	Y
PV22xS	N	N	N	N	N	N	N	N	Y	Y	Y	Y	Y	Y

4. PV22xS Enclosure Overview

The PV22xS SCSI Enclosure provides three configuration topologies, also referred to as modes, which are end-user configurable: [Joined](#), [Split](#), and [Cluster](#) topologies, as shown in Figure 4-1. A 3-position switch located on the Split Bus Module (SBM) sets the desired enclosure topology. Each of these topologies will be further explained in more detail below. The configuration will not take effect until the enclosure is power cycled which is designed to protect from accidental topology changes. SCSI Enclosure Services (SES) management features are an integral part of the enclosure. SES provides a unique set of features to monitor critical environmental variables such as temperature and power status.

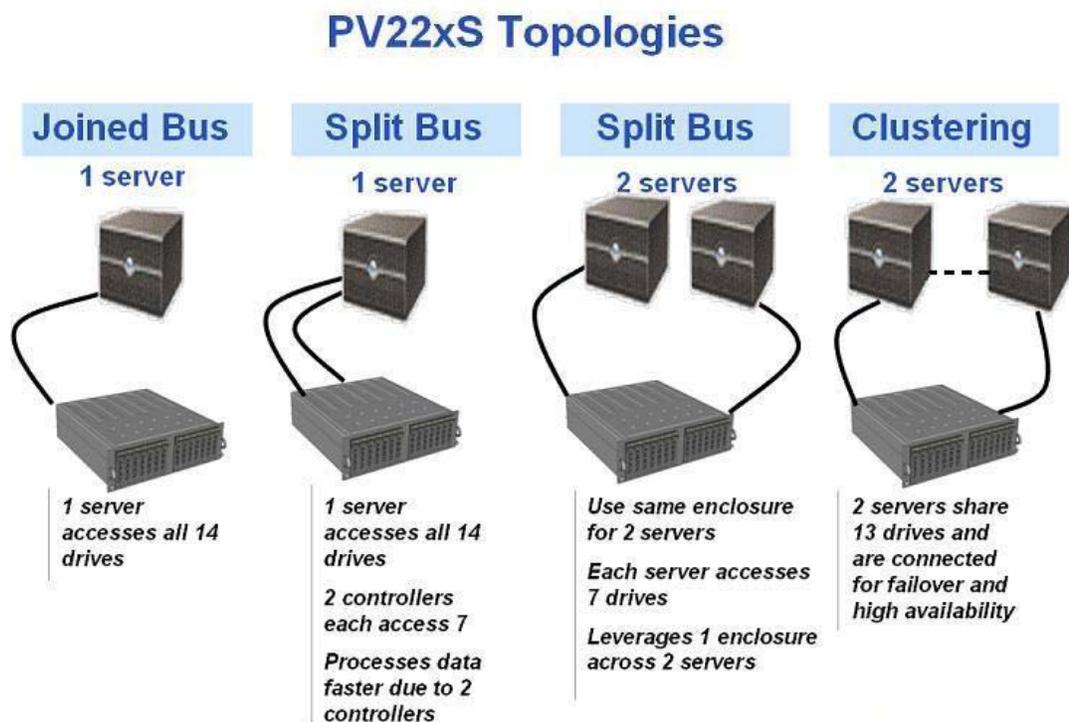


Figure 4-1 PV22xS Topologies

Joined Bus Topology

This is the simplest topology of operation providing access to all 14 hard drives on one shared SCSI bus. Only one cable must be attached between the server and the primary Enclosure Management Module (EMM) in the enclosure. If operating with two EMM modules, which provide redundant SES services, a second cable should not be attached to the secondary EMM. There can only be one physical path from the enclosure to the server. Attaching a second cable is not supported and may lead to SCSI bus address contention. If SES redundancy is not required, a Terminator module can be substituted in place of secondary EMM.

The primary EMM must always be present. Regardless of whether SES redundancy is available, the secondary EMM slot must always have either a EMM or Terminator module installed to properly terminate the SCSI bus. Leaving this slot unpopulated is not supported and will have undesirable consequences.

Configurations with two EMMs and one cable connected provide a level of redundancy. Even though there is no data path redundancy, there is enclosure management redundancy. The redundancy, however, is provided in the form of the SES enclosure management services to monitor and control critical functions such as temperature, fan speeds, initiating warnings/alerts, and in the worst case scenarios to safely power down the enclosure to prevent any potential damage under adverse thermal conditions. SES functions by default are performed by primary EMM. In the event there is a catastrophic failure of the primary EMM, the secondary EMM automatically will detect this event and perform enclosure management functions.

Split Bus Topology

This topology splits the PV22xS into two halves, essentially providing two independent, electrically isolated, SCSI busses capable of supporting up to seven (7) HDD on each bus. Because the two busses are electrically isolated, this allows simultaneous SCSI transactions on each bus. Two cables are required, one connected to each EMM. The cables can attach to two host systems (Figure 4-2) running different applications and operate completely independent from one another, or they can be connected to two SCSI channels on a single host system (Figure 4-3) to increase performance by splitting the I/O traffic between two different physical data paths. In either case, it is highly recommended that both halves of the PV22xS be connected to the same type of SCSI RAID controllers (e.g. PERC4 and PERC4 or PERC3 and PERC3). The primary EMM monitors the state of the seven HDD in slots labeled 0 thru 5, and 8. The secondary EMM monitors the state of the seven HDD in slots labeled 9 thru 15. SES management is provided by both EMM modules to their respective SCSI hosts. SCSI failures on one bus will not affect the other bus. Power and enclosure cooling, however, are shared. Power disruptions (e.g., power short on a HDD) will affect both busses equally.

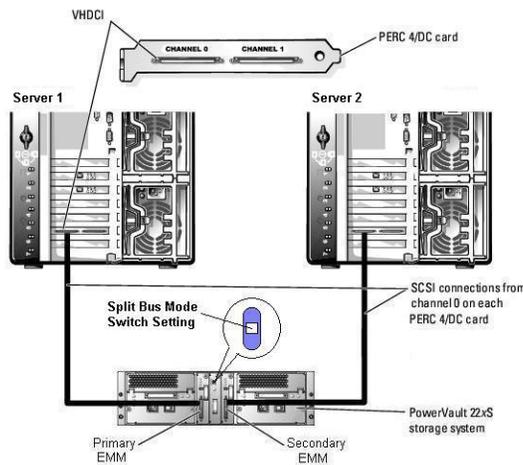


Figure 4-2 Split Bus Cabling to Two Host System

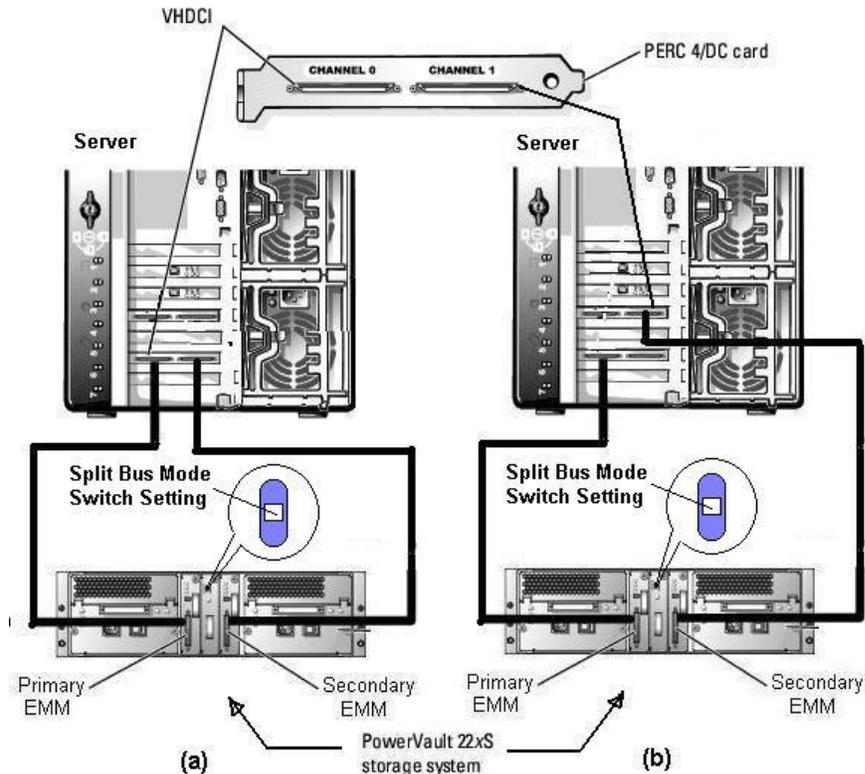


Figure 4-3 Split Bus Cabling Options to a Single Host System, (a) Same Adapter, (b) Different Adapters

Cluster Topology

Cluster Topology is very similar to Joined bus topology. It provides the capability to have a single shared SCSI bus with the additional capability of allowing concurrent access from two host controllers (Figure 4-4). In this topology, only 13 HDD are available since the PERC controller on the second host system uses up one of the available SCSI addresses (refer to Figure 4-5 for more details on the specific SCSI ID reassignments).

[Dell™ High Availability Clustering](#) (for more information refer to Appendix – A: References) provides the capability of configuring the PV22xS Enclosure to attach simultaneously to two different host systems. The hardware configuration of both host systems, including the SCSI RAID controller, must be identical. This is ideal for High Availability scenarios where access to the same data from two different systems is desirable. Clustering falls into two categories: Active/Active and Active/Passive.

Microsoft® Cluster Service MSCS and all PowerEdge™ Clusters support both active/active and active/passive cluster configurations. The term active/active refers to a cluster with at least one virtual server running on each node. For example, in a 2-node cluster configuration, when an application is running on Node 1, Node 2 does not need to remain idle waiting for Node 1 to fail. Node 2 can run its own cluster-aware applications (or another instance of the same application) while providing failover capabilities for resources on Node 1. Node 1 in turn can provide failover capabilities for resources on Node 2. Although both nodes have physical access to the same storage, the disk arrays are mutually exclusive and cannot be simultaneously accessed by both nodes. In the event of

a failover condition, access rights to the disk array assigned to the failed node is transferred to the other node. An active/active cluster node must be sized appropriately to handle the load of both nodes (in the event of a failover).

The term active/passive refers to failover cluster configurations in which one or more cluster nodes is actively processing requests for cluster applications while at least one cluster node simply waits for an active node to fail. An active/passive configuration is more costly in terms of price/performance because one or more servers sits idle most of the time. It is appropriate for business-critical systems since the application can use the full power of another server in case of a failure.

Dell PowerEdge Clustered systems fully support both cluster topologies with MSCS.

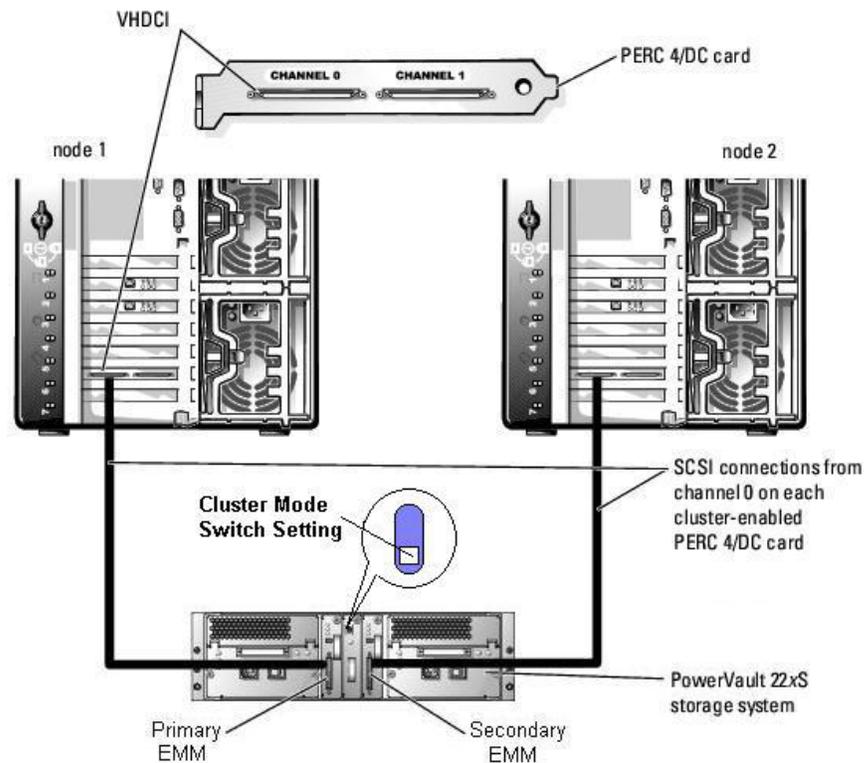


Figure 4-4 Cluster Topology Cabling Example

PV22xS – Common Features and Capabilities

Regardless of what topology is ultimately selected, the PV22xS provides a basic set of common features and capabilities.

SCSI ID Assignment

SCSI standard requires that all devices on the SCSI bus have a unique address, called a SCSI ID, with a maximum of 16 devices that can reside on the bus. The SCSI ID for the enclosure devices are hardwired by the enclosure backplane. PV22xS supports 15 of these devices with the SCSI controller in the host system using the remaining SCSI ID. In Cluster topology, however, the enclosure will give up one of its HDD SCSI ID, which will

be reserved for use by the second SCSI controller. A mapping of the SCSI ID address space reserved by the PV22xS enclosure is shown in Figure 4-5. The HDD identifiers are directly mapped to it's assigned SCSI ID, as shown in Figure 4-6.

Configuration	Cables Used	SCSI IDs Used																	
		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
Joined-bus	1																	H	S
Split-bus—primary EMM	1																	H	S
Split-bus—secondary EMM	1																	H	S
Cluster	2		S															H	H

NOTE: The unshaded SCSI IDs are available for hard-drive use as indicated for each configuration. The reserved SCSI IDs are used as follows:
 H = used by the host system initiator.
 S = used by the storage system SES.

Figure 4-5 SCSI ID Assignments

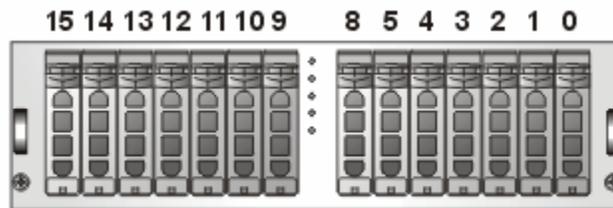


Figure 4-6 SCSI ID Numbers and Associated Hard Disk Drives

Multiple PV22xS Enclosures on a Single Host System

Daisy chaining of Multiple PV22xS enclosures is not supported. This is due to SCSI limitation of only allowing a maximum of 16 devices on one SCSI bus. Daisy chaining would effectively concatenate the devices in two (or more) enclosures into a single SCSI bus, thus violating the unique SCSI ID requirement.

Multiple PV22xS enclosures, however, can still be configured for a single host. The enclosure must be cabled directly to a separate SCSI channel on that host. A SCSI channel does not have to be on a physically separate adapter. For instance, a multi-channel SCSI controller such as a PERC4/DC is considered to be two SCSI channels even though both channels reside on one physical PCI adapter. In this example, two PV22xS can be attached to the PERC4/DC adapter. Conversely, it is also perfectly acceptable to use two PERC4/DC adapters and attach an enclosure to each. The maximum number of enclosures that can be configured for a single host system is limited by the number of channels the host system can support.

Mixing EMM Speeds

There are two types of EMM modules that PV22xS supports: U160 and U320. The numbers refer to the SCSI bus speed (in MB/sec) that the EMM is capable of supporting. For example, U160 can transfer up to 160MB/sec while U320 EMM supports up to 320MB/sec transfer rates. One or two EMM modules may be installed in the enclosure. When two are installed, it is required that they have the same transfer speed, which can be easily identified with labels on back of EMMs. Mixing of U160 and U320 EMM

modules is not supported. For more details, refer to section 5 (SCSI Host Bus Adapters / SCSI RAID).

Actual transfer rate of the SCSI bus is determined by host channel, EMMs and HDDs installed in PV22xS.

Mixing HDD Capacities

Mixing of different drive sizes is supported, but not recommended. PV22xS currently supports 18GB, 36GB, 73GB, 146GB, and 300GB capacities, with future support for higher capacities as they become available. If operating in a RAID configuration, the lowest capacity drive may define the RAID virtual disk size. For example, with two different HDDs, one 36GB and one 146GB drive, a RAID1 size will be limited to 36GB. On a RAID0 array size will be limited to 72GB. The excess space that is not included in the RAID array may still be partitioned and used. For more details on RAID array configurations, refer to section 5 (SCSI Host Bus Adapters / SCSI RAID).

Mixing HDD Spindle Speeds

Mixing of different drive spindle speeds is supported; however, it is not recommended. PV22xS currently supports 10K and 15K rpm spindle speeds, with future support for higher spindle rates as they become available. If operating in a RAID configuration, the lowest spindle speed may dictate the maximum performance of the RAID group. For example, in a RAID5 configuration with 5 drives having 4 HDDs running at 15K RPM and one running at 10K RPM, effective data transfer rate may be limited to the performance provided by the 10K RPM drive since all HDDs in the array perform as a single entity.

SCSI Enclosure Services (SES)

Environmental parameters such as temperature and voltage are continuously monitored by the EMM module. The PV22xS works in conjunction with Dell OpenManage™ to monitor and manage these parameters. Trigger thresholds can be set to warn the user when these parameters exceed safe operating levels, and to take specific action to power down the enclosure when they exceed a critical threshold level for maximum safety and avoid potential damages before they occur. Both local visual indicators and audio alerts are provided by the enclosure. By default, the audio alerts on the enclosure are disabled from the factory and must be enabled by the customer using Dell OpenManage, if desired. In summary, the following SES functions are provided by the PV22xS enclosure:

- Report enclosure shelf faults
- HDD 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)

Warm and Hot Pluggable Features

The PV22xS enclosure is designed for warm and hot plugging to help prevent damage when pulling or replacing certain components. Warm and hot plugging are very similar concepts allowing for the safe removal and replacement of components while the power is still on. The main difference is that warm plugging requires any I/O operations to be

properly quiescent before the remove/replacement operation takes place, whereas hot plugging can occur at any time without prior notification. Hot plugging allows the user applications and disk activities to continue uninterrupted while maintenance or repair actions are taking place. Warm plugging requires synchronization with the software applications to temporarily halt disk I/O activity before the hardware is removed/replaced.

- Warm Pluggable Devices
 - EMM module
 - SBM module
 - Terminator module
- Hot Pluggable Devices
 - HDD (when disk I/O is still occurring, all activities must be halted first)
 - Power supply module
 - Fan modules
- Neither Warm nor Hot Pluggable Devices

The backplane is critical for proper enclosure operations and cannot be warm or hot plugged. Power MUST be removed from the system.

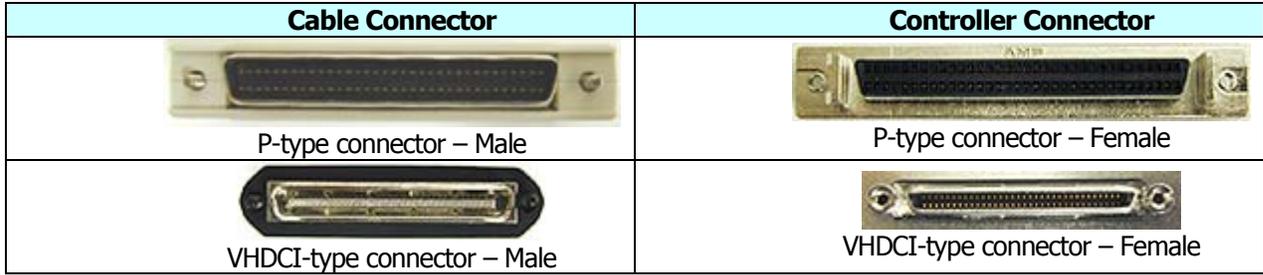
External SCSI Cables

Depending on the enclosure topology, one or two external SCSI cables can be attached between the host system(s) and the PV22xS enclosure. Although cables may physically look and feel the same, there are actually differences in the cable material and construction methods used that produces different quality levels of cable. In general, the higher the SCSI speed, the higher the quality of the cable that is required. SCSI cable quality is governed by the American National Standards Institute, Inc. (ANSI).

Specifically, for SCSI U160 speeds, a cable must comply with the ANSI SCSI Parallel Interface-3 (SPI-3) specification. For U320 speeds, ANSI SPI-4 must be adhered to. All PV22xS SCSI cables purchased from Dell have been tested and comply with the ANSI SPI-4 standard. When using existing customer cables or if the cables are purchased from 3rd party sources, ensure the cables are approved for ANSI SPI-4 operation. Using lower grade cabling will cause unpredictable results which can compromise the stability and reliable operation of the PV22xS enclosure.

External SCSI Connector Types

There are two types of SCSI connectors used by the external SCSI cables. One is called the High Density (HD) connector, also commonly referred to as a P-connector, and the other is the Very High Density Cable Interconnect (VHDCI). Both are 68-pin connectors, but the VHDCI has a smaller form factor, as shown in Figure 4-7. Two types of cables are available for the PV22xS, an HD-to-HD and HD-to-VHDCI. The EMM module uses the HD style connector where one end of the cable is attached. The other end will have either an HD or VHDCI connector. PERC3 and PERC4 series controllers use a VHDCI connector. Older controller cards generally use the HD connector. Always verify the controller connector type to ensure cables will be compatible.



• Figure 4-7 External SCSI Connector

PV22xS Configuration Topology and Application Attributes

PV22xS provides different configuration topologies, which may impact the storage solution for an application. All PV22xS configuration modes have various benefits for application attributes and are summarized as depicted in Table 4-1 and Table 4-2:

• Table 4-1: PV22xS Configuration Topologies Vs Application Attributes

Application Attributes	PV22xS Configuration Topologies		
	<i>Joined</i>	<i>Split</i>	<i>Cluster</i>
<i>Reliability</i>	Default	Default	Default
<i>Availability</i>	Default	Relatively better than Joined topology – depending on RAID level (see Table 4-2 for details)	High
<i>Serviceability</i>	All interconnect components are Field Replaceable units		
<i>Redundancy</i>	<ul style="list-style-type: none"> • Not on SCSI Path • Management capabilities with both ZEMMs 	<ul style="list-style-type: none"> • Not on SCSI Path • Management capabilities with both ZEMMs 	Provides both data and management
<i>Performance</i>	Depends on HDD and ZEMM technology – U160/U320		Relatively lower performance due to RAID controller configuration requirements
<i>Data Protection</i>	Default	Default	Default
<i>Scalability</i>	Maximum	Minimum (for a given topology)	Same as Joined, but one HDD less
<i>Cost</i>	Could be lower with Single ZEMM and non-redundant PSU	Requires two ZEMMs, but can be lower with non-redundant PSU	

• Table 4-2: PV22xS Configuration Topologies Vs RAID Level Performance

RAID Levels	PV22xS Configuration Topologies		
	<i>Joined</i>	<i>Split</i>	<i>Cluster</i>
<i>RAID 0</i>	Default RAID 0 performance	Best performance, if configured across multiple channels	Similar to Joined topology
<i>RAID 1</i>	Default	Increased Availability, if mirror set is on separate channel	Default
<i>RAID 5</i>	Default	Default	Default
<i>RAID 10</i>	Default	Higher performance than RAID 1 with same increased availability	Default
<i>RAID 50</i>	Default	Increase availability, if spanned across separate channels	Default

It can be observed that selection of a PV22xS configuration topology impacts not only Application Attributes but also impacts RAID level performance. Besides these configuration topologies, PV22xS provides some common capabilities as listed below for the storage solution:

Redundancy Capabilities

- *Physical SCSI Data Path Redundancy*

PV22xS using two Enclosure Management Modules (EMM) provides redundant data path when configured in Cluster topology.

- *Management Redundancy*

With multiple EMMs in PV22xS, hardware redundancy is provided when servicing enclosure management requests

- *Power Supply Redundancy*

PV22xS is capable of running on single power supply; multiple power supplies provide 1+1 redundancy

- *Cooling Redundancy*

PV22xS provides support of up to 2 fan units; each fan unit has 2 blowers in all making 4 fan blowers. The unit is thermally designed to continue operating with 1 of 4 fan blower failure, thus giving a 3+1 cooling redundancy

Reliability Capabilities

- *Thermal Reliability*

PV22xS consists of four digital temperature probes, which in conjunction with SES processor provides device protection against adverse thermal conditions. PV22xS is designed to operate at ambient temperature of 10°C to 35°C. Once the temperature probe detects ambient temperature to be below 0°C or above 50°C, SES processor notifies host of critical temperature range, which would then initiate a thermal shut down process. In case, PV22xS temperature probes detect ambient temperature either to be at or below -5°C or above 55°C, SES processor initiates a self shutdown process and hence protecting the device against adverse thermal conditions.

Serviceability Capabilities

- All components including backplane easily accessible and replaceable in rack
- LED indicators for visual status reporting on internal components
- Audible alarm to notify critical/non-critical status (for details please refer to PV22xS user's guide). Status is also reported to system management tools – OpenManage – OMSM or Array manager (for details refer to OpenManage user's guide)

5. SCSI Host Bus Adapters / SCSI RAID

This section addresses concerns specific to RAID. A summarized table (Table 5-2) is presented at the end of the section which provides an overview of the Dell SCSI and RAID solutions that support PV22xS.

Although attachment of PV22xS to 39160 SCSI HBA is supported, it is not recommended. 39160 is a U160 SCSI HBA, which is primarily intended for SCSI Tape devices, and would limit the performance of SCSI subsystem to U160 speeds. The controller doesn't provide any hardware RAID benefits similar to software RAID.

RAID: Hardware Vs Software

A RAID system mostly provides a fail-safe with the capability of improved performance data storage on a group of hard disk drives. RAID, when implemented with a hardware engine – including a dedicated RAID processor and memory system – is termed a hardware RAID. Whereas, when the RAID engine is implemented on a host processor and memory – either as part of operating system or driver, it is termed a software RAID.

Hardware and Software RAID implementations are associated with their respective pros and cons. The fact that the Hardware RAID is implemented using dedicated hardware resources gives a performance boost to the host systems which are required to perform heavy data processing. Software RAID performance scales based on the available memory and processing resources for a given host system. Thus, host systems which are used for heavy data processing will exhibit slower performance with Software RAID when compared to Hardware RAID implementations.

Selection of Hardware RAID vs Software RAID is a trade-off between performance and cost. Enterprise host systems which perform heavy data processing and are mission critical would be more suited for a Hardware RAID implementation. For more details on Hardware and Software RAID please refer to “Comparing RAID Implementation Methods” and “Implementing Software RAID on Dell PowerEdge Servers” papers in Appendix – A: References.

Considerations to Implement a RAID Solution

Implementing a RAID solution requires that certain considerations be addressed in order to achieve an optimum solution. These considerations are impacted by:

- Controller Family
- Data Protection Strategy
- High Availability Needs - Clustering
- Heterogeneous vs Homogeneous Environment – SCSI bus technology (U160/U320) and Hard Disk Drive (HDD) spindle speeds
- Interface Channel options - Internal and External channels
 - Spanning across channels/controllers/PV22xS
- Capacity Needs – Virtual Disk Size

Each of these above factors is discussed in detail in following sections.

Controller Family

Aside from the technical, physical and feature differences between the controller families (see Table 5-2), there are a few critical differences which impact behavior of these controllers and hence a storage solution. Dell's Storage controllers can be broadly categorized into two families – Adaptec PERCs and LSI PERCs. Design of these families impact the customer experience and may lead to ambiguities when configuring more complex storage strategies/solutions.

Though most of these differences are transparent to the user, when using the recommended Dell OpenManage™ software to configure storage subsystem; however, in cases where the BIOS utility is the preferred tool, an in-depth understanding of controller behavior is required to avoid possible issues. These differences include:

- *Controller BIOS utility terminology and operation*
 - The user interfaces of the boot-time BIOS utilities for the different controller families differ in layout and terminology. In environments where controllers from both families are being managed via the BIOS utilities, these differences can lead to confusion. Refer to user guides for controller specific information on the functionality and terminology used by these utilities.

NOTE: In the context of this document, the word 'array' (LSI BIOS) = 'container' (Adaptec BIOS) = 'logical drive' (LSI POST) = 'volume' (OpenManage Software).
- *Layout of multiple logical arrays on the same group of disks*
 - LSI family of controllers require that once an array is created on a group of disks, any subsequent arrays created on any of these same disks can only be created on the same group of disks; subsequent arrays are not allowed to use any disks not in the original group, nor any subset of the original group. Adaptec controllers are less restrictive in this regard.
 - Similarly LSI controllers are more restrictive of hot spare functionality in this scenario. If redundant and non-redundant arrays are placed on the same group of disks and a hot spare is assigned to the redundant array, because of the non-redundant array's failed state, the hot spare will never start the rebuild on the redundant array.
- *Mixed-size hard drives within a RAID10/RAID50*
 - RAID50 not supported on Adaptec controllers. This configuration is supported on LSI controllers only
 - When mixing drive sizes in RAID10/RAID50 on LSI controllers, with default maximum size, the controller will not coerce the larger 'child' arrays to match the size of the smallest 'child' array. Controller will stripe across all 'child' arrays until the smallest one is full, then it will continue to stripe across the remaining 'child' arrays until the next one is full, and so on until all the available space is used. This can impact performance as the array fills up, since the number of stripes decreases as 'child' arrays fill up.

In contrast, Adaptec controllers will only use space on all 'child' arrays equal to the available space on the smallest 'child' array, and will leave the remaining space on the larger drives unused and available, and the performance behavior will not be affected as the array fills up.

- *Array Reconfiguration / Disk Expansion / Rebuild*
 - As stated in the 'Getting Started with RAID' guide (see Appendix A – References), the recommended method of changing an array configuration is to back up the data on the array, create a new array with the new desired parameters, and then migrate the data to the new array. In cases where this is not possible, reconfiguring arrays is supported for the controllers. However, should an array need to be reconfigured, the reconfiguration paths supported differ on different controllers. Please refer to the documentation on each controller for information on supported reconfiguration paths.
 - Many factors may influence the time it takes to perform a rebuild or reconfiguration operation on an array including: Size of disk and array, RAID level, IO profile, system load, PERC family,

For more details on the rebuild process please refer to Appendix A – Reference 5: Maintenance Best Practices for Direct Attached SCSI Solutions

- *Controller upgrade / Array Migration / Drive Roaming / metadata*
 - PERC controllers store array configuration metadata both in controller NVRAM and on-disk. Arrays are upgradeable and can be migrated within the family of controllers made by the same vendor; i.e., migration of arrays from one LSI controller to another or from one Adaptec controller to another is possible whereas migration between the two families is not; e.g., an array from an LSI controller cannot be migrated to an Adaptec controller, or vice-versa.
 - Drive roaming is initiated when the HDDs are changed to different SCSI IDs or channels on a same controller. When the HDDs are placed on different SCSI IDs or different channels of the same controller, controller detects the RAID configuration from configuration metadata residing on the HDD. Drive roaming is supported on all families of SCSI RAID controllers.

NOTE: In a clustering environment, drive roaming is supported within the same channel only.

- For additional information on Array Migration or Drive Roaming, please refer to the User Guide for the controller.
- *Write cache policy considerations*
 - Write back

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, but on PERC controllers that do not have a battery

back up unit (see Table 5-2), there is a risk of data loss if system power is lost while unwritten information is still in the cache.

- Write Through

In write-through caching, a data transfer is completed when the disk subsystem receives all of the data from the host. There is no risk of data loss in a power-loss scenario (as in write back), but it is slower, since the controller must wait for the hard drive to return a good status on the write to the controller before proceeding to the next operation.

If the PERC controller has a battery backup unit (see Table 5-2) or if a host system is powered by an alternate power source (e.g. UPS etc), the controller's cache retains data during a power loss. Thus PERCs with battery back-up not only provide the performance benefit of write-back caching but also help ensure integrity of data during system power loss. For data retention time please refer to Table 5-2.

Data Protection Strategy

- Battery back up: see "Write Cache policy considerations" section above.
- Global hot spare
 - A global hot spare can be used for any array that is on the same controller as the hot spare.
 - More cost-effective than dedicated hot spares, but does not protect multiple arrays as completely as dedicated hot spares.
 - A global hot spare is selected randomly and may be in a different enclosure than the failed drive it is replacing. Can impact performance if rebuilt array spans enclosures/channels.
- Dedicated hot spare
 - A dedicated hot spare is assigned to one or more arrays.
 - Dedicated hot spares that reside in the same storage enclosure typically have better performance than global hot spares.
 - Dell recommends using dedicated hot spares for critical data.
- Consistency checks
 - Consistency checks verify the correctness of redundant data in an array.
 - As during a rebuild, array performance will be degraded while check is running.
 - Some PERC controllers allow pausing of consistency check and resume it later or to resume the consistency check after the system reboots.
 - Recommended as part of regular system maintenance for critical data.

For more details on rebuild process please refer to Appendix A – Reference 5: Maintenance Best Practices for Direct Attached SCSI Solutions

- **Patrol Read**
 - Patrol Read is a preventative maintenance background operation (available only on PERC 3 (except PERC 3/DI), PERC 4 and PERC 4e controller families running 3.0 and higher firmware).
 - Patrol Read examines each block in a configured physical disk for media errors and attempts to fix media errors by reallocating the bad block.
 - Patrol Read is designed to run so as to not impact I/O performance.
 - Patrol Read has two run modes: Manual and Automatic. The mode is set from within the controller BIOS. By default the run mode is set for “Automatic”.
 - Highly recommended as part of regularly scheduled system maintenance to prevent downtime due to media errors.

NOTE: Patrol Read and Consistency Check are not the same operations; Patrol Read examines the physical disk for media defects while Consistency Check validates data using data parity, to help prevent data errors.

High Availability Needs – Clustering

- Clustering is supported only on some HBAs – see Table 5-2. Clustering is not supported on embedded or non-RAID controllers.
- Multiple logical volumes on the same physical array are not supported in clustering environment – only one logical volume/array.

Heterogeneous vs. Homogeneous environments

- SCSI Technology (U320/U160/Ultra 2)

It is possible to mix a combination of hard drives compliant with U320, U160 or Ultra-2 generation of SCSI technologies. Mixing of hard drives is supported as long as all the devices on the SCSI bus have Low Voltage Differential (LVD) signaling.

Though possible, it is not recommended to mix drives, since overall SCSI bus performance will be limited by the slowest technology present in the mix of array. Where mixing of HDD technology is not avoidable, it is a best practice to have all hard drives in any particular array be of matching technology. In all cases, the overall subsystem performance will be limited to some degree by the slowest technology present in the mix.

- Spindle Speed

It is recommended that all hard drives within any particular array be of the same spindle speed (RPM – which can vary from 10,000-15,000 on the hard

drives supported by the PV22xS). Mixing spindle speeds within the same array will impact array performance, especially in a striped (RAID0, RAID10 or RAID50) environment. Higher RPM drives will have to wait for the slower ones to complete the operations. This is the case regardless of cache policy settings.

Interface Channel Options – Internal/External/Multi channel device configuration

- See Table 5-2 for internal/external connector details for each controller.
- Do not connect devices both internally and externally on any one channel.
- Termination
 - The PV22xS is also already terminated. No external SCSI bus terminators are needed when connecting to a PV22xS.
- Spanning of Arrays

Dell's PERC controllers provide additional array configuration options of spanning logical arrays. Dell supports spanning of logical arrays only across multiple channels on same PERC controllers. If multiple PV22xS are configured (Joined or Split topology - Figure 4-3a) on the same PERC then a logical array may be spanned across multiple PV22xS.

In general, spanning of logical arrays across multiple channels on a controller improves data throughput. With the introduction of another data path, depending on the RAID level, the configuration may compromise data availability. When configuring logical arrays across multiple channels it is essential to understand the trade-off between performance, reliability and availability of data. Depending on the RAID level for an array, spanning may offer a better solution than without spanning. The following table lists pros and cons of spanning in relation to RAID levels:

• Table 5-1: Pros and Cons for Spanning Arrays with various RAID Levels

Level	Pros	Cons
<i>RAID 0</i>	Improves data throughput by optimizing bandwidth utilization	Reduces reliability and availability of data with additional point of failure – second channel
<i>RAID 1</i>	Improves data availability against channel failures, if mirrored set is spanned across separate channel	If PV22xS is used on both channels (either in Split or Joined topology), limits the usage of PERC on internal channel (connecting devices to the same channel on internal and external connectors is not supported)
<i>RAID 5</i>	May improve data throughput for certain accesses, if data and parity for the access span across channels	Even though RAID5 provides data protection but compromises reliability and availability of data at SCSI channel level
<i>RAID 10</i>	Improves data availability against channel failures, if mirrored set is spanned across separate channel	If PV22xS is used on both channels (either in Split or Joined topology), limits the usage of PERC on internal channel (connecting devices to the same channel on internal and external connectors is not supported)
<i>RAID 50</i>	May improve data throughput, similar to spanned RAID5	Reduces reliability and availability of data as compared to non-spanned RAID50 at SCSI channel level

Note: Dell PERC controllers don't support spanning of arrays across different controllers. Spanning across channels on the same PERC controller is only supported.

Capacity Needs – Virtual Disk Size

Virtual disks on SCSI PERC controllers cannot be created with an aggregate size greater than 2TB due to limitations on SCSI controllers and operating systems. Host management application (OMSM, Array Manager) will give an error if an attempt is made to make virtual disk greater than 2TB, prompting user to reduce size of virtual disk. For more information please refer to “Beyond The 2-TB SCSI Logical Unit” white paper at:

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

Additional RAID information

- Dell ‘Getting Started with RAID’ document

<http://support.dell.com/support/edocs/storage/RAID/RAIDbk0.pdf>

Table 5-2 – Controller Families & Features: Part 1

	PERC 4e/Di	PERC 4e/DC	PERC 4/Di (PE2600)	PERC 4/Di (PE1750)
SCSI Technology	Wide Ultra320	Wide Ultra320	Wide Ultra320	Wide Ultra320
Host Bus Interface	x4 PCI Express (Embedded)	X8 PCI Express (HBA)	PCI 2.2 (Embedded)	32-bit / 33 MHz PCI (Embedded)
# Channels	2	2	2	2
# internal connectors	2 x 68-pin HD	2 x 68-pin HD	2 x 68-pin HD	n/a (hardwired to backplane)
# External Connectors	1 x 68-pin UHD	2 x 68-pin UHD	1 x 68-pin HD	1 x 68-pin UHD
RAID Levels Supported	0, 1, 5, 10, 50	0, 1, 5, 10, 50	0, 1, 5, 10, 50	0, 1, 5, 10, 50
Cluster Support?	no	yes	no	no
Cache Size (DIMM)	256MB DDR2	128MB DDR	128MB DIMM	128MB DDR2
Write Cache Options	off, write-through, write-back	off, write-through, write-back	off, write-through, write-back	off, write-through, write-back
Read Cache Options	read-ahead (default), off (selectable)			
BBU Battery Life	48-72 hours	48-72 hours	48-72 hours	48-72 hours
Max # Logical Drives/Controller	40	40	40	40
Max # Physical Drives/Logical Drive	32	32	32	32
Vendor	LSI	LSI	LSI	LSI
SCSI Chip	LSI 53C1030	LSI 53C1030	LSI 53C1030	LSI 53C1030
RAID Processor	Intel GC80332	Intel GC80332	Intel GC80303	Intel i80321

Table 5-2 – Controller Families and Features: Part 2

	PERC 4/SC	PERC 4/DC	PERC 3/SC	PERC 3/DCL
SCSI Technology	Wide Ultra320	Wide Ultra320	Wide Ultra3 / U160	Wide Ultra3 / U160
Host Bus Interface	64-bit / 66 MHz PCI 2.2 (HBA)			
# Channels	1	2	1	2
# internal connectors	2 x 68-pin HD	2 x 68-pin HD	1 x 68-pin HD	2 x 68-pin HD
# External Connectors	1 x 68-pin UHD	2 x 68-pin UHD	1 x 68-pin UHD	2 x 68-pin UHD
RAID Levels Supported	0, 1, 5, 10, 50	0, 1, 5, 10, 50	0, 1, 5, 10, 50	0, 1, 5, 10, 50
Cluster Support?	no	yes	no	No
Cache Size (DIMM)	64 MB	128MB	64 or 128MB	64 or 128MB
Write Cache Options	off, write-through, write-back	off, write-through, write-back	off, write-through, write-back	off, write-through, write-back
Read Cache Options	read-ahead (default), off (selectable)			
BBU Battery Life	no BBU	48-72 hours	no BBU	no BBU
Max # Logical Drives/Controller	40	40	40	40
Max # Physical Drives/Logical Drive	32	32	32	32
Vendor	LSI	LSI	LSI (AMI)	LSI (AMI)
SCSI Chip	LSI 53C1020	LSI 53C1030	Qlogic 12160	Qlogic 12160
RAID Processor	Intel GC80302	Intel GC80303	Intel i960	Intel i960

Table 5-2 – Controller Families and Features: Part 3

	PERC 3/DC	PERC 3/QC	PERC 3/Di	PERC 2/QC	Adaptec 39160
SCSI Technology	Wide Ultra3 / U160	Wide Ultra3 / U160	Wide Ultra3 / U160	Wide Ultra 2	Wide Ultra3 / U160
Host Bus Interface	64-bit / 66 MHz PCI 2.2 (HBA)	64-bit / 66 MHz PCI 2.2 (HBA)	64-bit / 66 MHz PCI 2.2 (Embedded)	64-bit / 33 MHz PCI 2.1 (HBA)	64-bit / 66 MHz PCI 2.2 (HBA)
# Channels	2	4	2	4	2
# internal connectors	2 x 68-pin HD	2 x 68-pin HD	2 x 68-pin HD (except 1650 & 2650 - N/A)	2 x 68-pin HD	2 x 68-pin HD (LVD); 1x 50-pin Ultra
# External Connectors	2 x 68-pin UHD	4 x 68-pin UHD	0, 1 or 2 x 68-pin HD (server platform dependant)	4 x 68-pin UHD	2 x 68-pin UHD (LVD)
RAID Levels Supported	0, 1, 5, 10, 50	0, 1, 5, 10, 50	0, 1, 5, 10	0, 1, 5, 10	n/a
Cluster Support?	yes	no	no	no	no
Cache Size	64 or 128MB DIMM	64 or 128MB DIMM	128MB DIMM	16 - 128 MB DIMM	n/a
Write Cache Options	off, write-through, write-back	off, write-through, write-back	off, write-through, write-back	off, write-through, write-back	n/a
Read Cache Options	read-ahead (default), off (selectable)	read-ahead (default), off (selectable)	read-ahead (default), off (selectable)	read-ahead (default), off (selectable)	n/a
BBU Battery Life	48-72 hours	48-72 hours	48-72 hours	48-72 hours	n/a
Max # Logical Drives/Controller	40	40	64	64	n/a
Max # Physical Drives/Logical Drive	32	32	64	64	n/a
Vendor	LSI (AMI)	LSI (AMI)	Adaptec	Adaptec	Adaptec
SCSI Chip	Qlogic 12160	2 X Qlogic 12160	Adaptec 7899	2 x Adaptec 7897N	Adaptec 7899
RAID Processor	Intel i960	Intel i960	Intel i960 OR GC80303 (server platform dependant)	StrongARM SA110S	n/a (non-RAID)

6. Summary and Conclusions

After analyzing and assessing the importance of all the storage components relative to attributes essential for performance of any given application, it can be observed that not only understanding the characteristics of each storage component is critical but also how the storage component is configured has direct impact on the effectiveness of an application.

Based on the study of each storage component relative to attributes essential for effectiveness of an application, Table 6-1 lists typical storage recommendations for various applications. Recommendations listed in the table are only an example for a storage solution; each actual storage solution will have additional factors which need to be considered before determining an optimum storage solution. These additional factors 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

• Table 6-1: Typical Storage Recommendation for Different Storage Applications

Application		Storage Recommendations			Remarks
<i>Servers</i>	<i>Key Parameters</i>	<i>RAID Level</i>	<i>PV22xS</i>	<i>Host Controller</i>	
Email	* Availability * Performance * Scalability	RAID10 Provides both Availability and performance	Cluster Topology Scalable and high availability	For added performance enable Write back cache, if available on PERC	Recommendation targets a typical Enterprise Email Server
Database	* Availability * Performance	RAID10 Provides highest availability of all other RAID levels	Cluster Topology Provides availability of all PV22xS topologies	If added performance is required enable Write back cache, if available on PERC	Performance requirements for DB servers depend on number of users
Web	* Performance	RAID0 in Web Server farms Provides maximum performance of all RAID levels	Split Topology Spanning RAID0 across channels provides best performance	PERC, with 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	RAID10 Provides highest availability of all other RAID levels	Cluster Topology Provides availability of all PV22xS topologies	For added performance enable Write back cache, if available on PERC and data protection at PERC level	Recommendation targets Enterprise applications which require higher availability than performance. In general for all OLT data protection is of utmost importance
File	* Scalability	RAID5 Maximizing storage with some data protection	Joined Topology Provides maximum scalability	PERC with Global Hot Spare will improve availability for RAID5	These servers usually have data backed up periodically and hence, require minimum data protection from Storage components
Streaming Media	* Performance	RAID0 Provides maximum performance of all RAID levels	Split Topology Spanning RAID0 across channels provides best performance	PERC, with read ahead capabilities for added performance boost	Performance is of utmost importance for Streaming media servers. Data on these servers is usually backed up on external medias such as Tape libraries or optical medias - CD/DVD

7. Appendix – A: References

1. [ANSI T10/1365D SCSI Parallel Interface – 4 \(SPI-4\) and SCSI Block Commands – 2 \(SBC-2\)](#)
(<http://www.t10.org>)
2. [Dell™ PowerVault™ 220S and 221S Systems Documents](#)
(<http://support.dell.com/support/edocs/stor-sys/spv22xs/en/index.htm>)
3. [Dell 'Getting Started with RAID' document](#)
(<http://support.dell.com/support/edocs/storage/RAID/RAIDbk0.pdf>)
4. [Dell High Availability Cluster](#)
(http://www1.us.dell.com/content/topics/global.aspx/solutions/en/clustering_ha?c=us&cs=555&l=en&s=biz)
5. [Maintenance Best Practices for Direct-Attached SCSI Solutions](#)
(<http://www.dell.com/downloads/global/products/pvaul/en/Maintenance%20Best%20Practices%20for%20Direct-Attached%20SCSI%20Solutions.pdf>)
6. [Beyond The 2-TB SCSI Logical Unit](#)
(http://www.dell.com/downloads/global/vectors/2004_2tblun.pdf)
7. [Comparing RAID Implementation Methods](#)
(http://www1.us.dell.com/content/topics/global.aspx/power/en/ps3q03_dumouchelle?c=us&cs=555&l=en&s=biz)
8. [Implementing Software RAID on Dell PowerEdge Servers](#)
(http://www1.us.dell.com/content/topics/global.aspx/power/en/ps3q03_nadon?c=us&cs=555&l=en&s=biz)

8. Appendix – B: Glossary

Array – Collection of multiple physical Hard Disk Drives to form a single logical volume

BBU – Battery Backup Unit

BIOS – Basic Input Output System

Cluster – Multi-initiator Environment with MSCS operating environment

HBA – Host Bus Adapter (non-RAID controllers)

HDD – Hard Disk Drive

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

PERC – Power Edge RAID Controller

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

ROMB – RAID On Motherboard

SCSI – Small Computer System Interface

U160 – Ultra 160 (160MB/sec throughput)

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

Warm Plug – Insertion or removal of device with all the I/Os quiescent

9. Appendix – C: About the Dell SCSI Storage Solution Team

Authors

David Papasan – International Product Support
Gary Mott – PowerVault™ 22xS Product Engineering Manager
John Woo – PowerVault 22xS Product Engineering Lead
Mickey McGee – RAID Engineer
John Langdon – RAID Engineer
Phil Curran – PowerVault 22xS Marketing
Ryan King – International Product Support
Samit Ashdhir – PowerVault 22xS Product Engineer

10. Appendix – D: Revision History

<i>Revision</i>	<i>Description</i>	<i>Date</i>	<i>Modified By</i>
A00	First Release	04/28/2005	SCSI Storage Solution Team
A01	Updated “™” and “®” information where applicable	05/06/2005	Samit Ashdhir
A02	Updated requirement for controllers to be attached to PV22xS in Split and Cluster Topologies (Section 4) Updated with brief on “Hardware Vs Software RAID” (Section 5) Updated “Drive Roaming” and “Patrol Read” (Section 5)	11/17/2005	SCSI Storage Solution Team