

BEST PRACTICES FOR MICROSOFT EXCHANGE SERVER 2007 AND POWERVAULT MD FAMILY

PowerVault MD3000 and
MD3000i



Table of Contents

INTRODUCTION	3
DISCLAIMER	3
BASICS OF PERFORMANCE TUNING	3
UNDERSTANDING THE CONTEXT FOR PERFORMANCE TUNING	4
THREE COMPONENTS THAT INFLUENCE PERFORMANCE	4
MICROSOFT EXCHANGE SERVER 2007 CONSIDERATIONS	4
ANALYZING I/O CHARACTERISTICS.....	5
CALCULATING THEORETICAL USAGE	5
STORAGE GROUP AND DATABASE DESIGN	7
Locating Database within Storage Groups.....	7
Backup and Restore Factors.....	7
DETERMINING THE BEST RAID LEVEL	7
CHOOSING THE NUMBER OF PHYSICAL DISKS TO PUT IN A DISK GROUP	8
CLUSTERING AND DATA PROTECTION	8
CONSIDERING THE SERVER PLATFORM	9
CONSIDERING THE SERVER HARDWARE ARCHITECTURE	9
Calculating Aggregate Bandwidth.....	9
Sharing Bandwidth with Multiple HBAs.....	9
HOST CONNECTIVITY	9
MICROSOFT WINDOWS SERVER CONSIDERATIONS	10
Aligning Host I/O with RAID Striping.....	10
NTFS Allocation Unit Size	11
SETTING UP THE STORAGE ARRAY	11
FACTORS INFLUENCING STORAGE PERFORMANCE.....	11
ESTIMATING CAPACITY LIMITS	11
CALCULATING OPTIMAL SEGMENT SIZE.....	11
SETTING THE GLOBAL MEDIA SCAN	12
SETTING LUN-SPECIFIC PARAMETERS	12
Setting the LUN-Specific Media Scan.....	12
Setting the LUN-Specific Read Cache and Read Prefetch	13
LOCATING DISK GROUPS.....	13
Striping Disk Groups in RAID Controller Modules.....	13
Cabling the Dell PowerVault MD3000 and MD3000i.....	13
USING THE POWERVAULT MODULAR DISK STORAGE MANAGER PERFORMANCE MONITOR	15
SUMMARY OF PERFORMANCE TUNING BEST PRACTICES	16
CONCLUSION	17
CONTACT INFORMATION	17
REFERENCES	18

Introduction

This document details the requirements necessary to implement a Dell MD3000 or MD3000i storage array in a Microsoft Exchange Server 2007 environment. This document is not intended to be an installation and configuration guide for Microsoft Exchange Server 2007, but rather a guide to tune and optimize the storage array within the Exchange environment.

For additional information about Exchange Server 2007, refer to the Microsoft Exchange server 2007 product page. Go to:

<http://www.microsoft.com/exchange/default.mspx>

Disclaimer

Due to the highly customizable nature of an Exchange environment, you must take into consideration the specific environment and equipment to achieve optimal performance from a Dell storage array. Also note that many procedures within this document should be attempted only by experienced Microsoft Exchange Server 2007 administrators and trained storage specialists with intimate knowledge of the working environment. This document results from the collaboration of numerous Dell engineers and performance specialists. Additional training can be obtained from Dell.

For more information on PowerVault MD3000i and MD3000 and Exchange check out:

Principled Technologies Performance Testing of PowerVault Storage and Exchange:

<http://www.dell.com/downloads/global/products/pvaul/en/exchange2007-compare-md3000i-hp.pdf>

PowerVault MD3000 ESRP Testing: <http://technet.microsoft.com/en-us/exchange/bb412164.aspx>

PowerVault MD3000i ESRP Testing:

http://www.dell.com/downloads/global/products/pvaul/en/dell_powervault_MD3000i.pdf

Basics of Performance Tuning

Understanding the Context for Performance Tuning

The challenge of storage performance tuning is to understand and control interacting factors while accurately measuring Microsoft Exchange Server 2007 performance. Because the performance of the storage array accounts for only a portion of the overall Exchange performance, tuning must be done in context. The full context includes other contributing factors such as the I/O characteristics of Exchange and the components in the data path. These contributing factors include the following:

- HBA
- Switches
- File system
- Operating system
- Server

With multiple parameters to consider, the task of performance tuning can seem difficult. To reduce the complexity of tuning, Dell's MD3000 and MD3000i storage arrays feature performance monitoring and flexible tuning controls in the PowerVault Modular Disk Storage Manager.

Three Components That Influence Performance

This document provides an overall approach to tuning I/O performance and also provides specific guidelines for using the storage array tuning controls. These recommendations begin with an overall analysis of the three elements that determine I/O performance:

- The Microsoft Exchange Server 2007 application
- Server platform (hardware, operating system, device drivers)
- Storage array

Microsoft Exchange Server 2007 Considerations

When planning for Exchange 2007 storage, performance should be the primary focus, with space planning a secondary concern. A system that is first optimized for Exchange I/O performance will generally result in adequate space as well.

Measuring or predicting the performance of transactional I/O facilitates in determining whether the configuration meets the requirements for acceptable disk latency and disk throughput. For Microsoft Exchange Server 2007, the storage array is considered to be performing poorly if average read and write disk latencies are greater than 20 ms for database virtual disks and 10 ms for log virtual disks. Frequent or extended latency spikes above 50 ms are also indicators of poor performance.

EXCHANGE AND POWERVAULT MD3000 FAMILY

It is critical that you develop an understanding of the Exchange Server 2007 I/O requirements for your environment. For additional information, refer to the Microsoft document, *What Causes Exchange Disk I/O?* Go to:

<http://technet.microsoft.com/en-us/library/bb738154.aspx>

Analyzing I/O Characteristics

Microsoft Exchange Server 2007 accesses data either randomly or sequentially, depending on the data being accessed.

Database access is in random 8 KB page sizes. However, I/O coalescing generates larger I/O sizes—up to 1 MB. By increasing the page size to 8 KB, an increase of 4 KB over Microsoft Exchange Server 2003, disk I/O decreases by allowing larger amounts of data to be accessed in a single I/O. Exchange 2007 databases also take advantage of the 64-bit architecture allowing for a larger database cache, which results in a decrease in the read percentage of the total disk I/O. A system configured in accordance with Microsoft's recommended memory configuration guidelines results in a database read-to-write ratio of approximately 1:1.

Log access is sequential with writes ranging from 512 bytes up to the 1MB buffer size. All changes made to the database are written to the transaction logs first. Exchange 2007 Server transaction log files are 1 MB in size, a decrease from 5 MB in Exchange Server 2003. Due to a decrease in database reads and a smaller log file size, the log-to-database write ratio is approximately 3:4.

Calculating Theoretical Usage

When planning a new Microsoft Exchange Server 2007 implementation, it is critical that you determine the total number of users in your configuration and the approximate user mailbox size before you create storage virtual disks. Without these values, you cannot accurately determine the amount of storage to allocate. Remember to allow for future growth in your calculations.

After you have determined the total number of users and their corresponding usage profiles, use the following procedure to predict the number of IOPS the Exchange configuration must support. [Table 1](#) shows the estimated user IOPS profile as well as recommended user mailbox size.

Table 1: User Profiles

User Type (Usage Profile)	Send/Receive Per Day (Approximately 50 KB Message Size)	Estimated IOPS Per User (Cached Mode)	Mailbox Size
Light	5 sent/20 received	0.11	100 MB
Average	10 sent/40 received	0.18	250 MB
Heavy	20 sent/80 received	0.32	1024 MB
Very Heavy	30 sent/120 received	0.48	2048 MB

EXCHANGE AND POWERVAULT MD3000 FAMILY

To determine the required database IOPS:

Number of users x (I/O profile of user + 20 percent for overhead) = required IOPS for database virtual disks

EXAMPLE

4000 (users) x 0.384 (heavy user profile + 20 percent overhead¹) = 1536 IOPS required for database virtual disks

To determine the required log IOPS:

Assume a database read-to-write ratio of 1:1. Continuing with the previous example, out of 1536 IOPS for the database virtual disks, there would be 768 IOPS for reads and 768 IOPS for writes.

Follow the log-to-database write ratio of 3:4 to determine the required IOPS for log virtual disks.

768 DB writes/sec x (.75) = 576 log writes/sec

To determine the total required IOPS:

Total required IOPS = required IOPS for database virtual disks + required IOPS for log virtual disks

EXAMPLE

1536 IOPS + 576 IOP = 2112 IOPS

Divide the total IOPS by the number of Exchange storage groups to determine the IOPS each storage group will need to support. When configuring the storage virtual disks, this value should be the minimum IOPS that the storage group virtual disk can support.

Total IOPS ÷ Number of storage groups = IOPS per storage group

NOTE Prior to creating storage array virtual disks it is highly recommended that you use the Exchange 2007 Mailbox Server Role Storage Requirements Calculator developed by Microsoft. This calculator helps you determine the storage requirements, including I/O performance and capacity requirements. This calculator also suggests the optimal LUN layout based on the input parameters.

For additional information about the Mailbox Storage Calculator for Exchange 2007, including usage instructions, refer to the Microsoft document, *Mailbox Server Storage Design*. Go to:

<http://technet.microsoft.com/en-us/library/bb738147.aspx>

¹Add 20 percent to the estimated IOPS to allow for additional Exchange activity during peak hours.

Storage Group and Database Design

Locating Database within Storage Groups

In Exchange Server 2007, Microsoft recommends locating all databases in a storage group on the same physical LUN. It is also a best practice to place no more than one database in each storage group. Do not place multiple storage groups on the same physical LUN: use only one storage group per LUN. Remember to verify that your chosen LUN design can accommodate the required IOPS per storage group.

Backup and Restore Factors

If you plan on taking advantage of the Dell VSS hardware provider to implement Exchange 2007 snapshot backups, design the LUNs to minimize impact to separate storage groups. Each Exchange 2007 storage group should contain a single transaction log LUN and a single database LUN that are independent of other storage groups. This design allows you to back up and restore a single Exchange storage group with no performance impact to other storage groups that might exist.

Determining the Best RAID Level

RAID 1/10² and RAID 5 are all viable options for Exchange 2007 database virtual disks because these RAID levels have similar characteristics for read environments. However, RAID 1/10 offers the best performance while still providing data protection by mirroring each physical disk. In an Exchange 2007 environment, the transaction logs are the most important data set and low write latency is critical for server performance. With that in mind, logs should be placed on RAID 1/10 virtual disks with battery-backed write cache. Latency for the transaction log virtual disks should be less than 10 ms.

Considering Exchange Server 2007's database read-to-write ratio of 1:1, RAID 5 virtual disks are at a disadvantage due to the increased writes. RAID 5 is challenged most by random writes, which can generate multiple physical disk I/Os for each host write. However, when designed properly, RAID 5 provides the required IOPS and does so using fewer disks than RAID 1/10. When you design your virtual disks, make sure to use enough disks in RAID 5 disk groups to support the projected IOPS requirement calculated earlier in the "Calculating Theoretical Usage" section on page 8. Latency for the database virtual disks should be less than 20 ms.

NOTE Different RAID levels can be tested by using the PowerVault Modular Disk Storage Manager Dynamic RAID Migration feature, which allows the RAID level of a disk group to be changed while maintaining continuous access to the data.

²PowerVault Modular Disk Storage Manager only shows the option RAID 1. However, if you select four or more drives, RAID 10 is automatically configured across the virtual disk—two drives for user data and two for the mirrored data.

Choosing the Number of Physical Disks to Put in a Disk Group

Microsoft Exchange Server 2007 is an I/O intensive application. Therefore, you must select a number of physical disks that matches the per disk group I/O rate needed to support the application. (See "Calculating Theoretical Usage" on page 5.) Make sure to account for the I/Os required to implement the data protection of the selected RAID level. You can add more physical disks to the configuration for a linear increase in performance, up to the point of RAID controller module saturation. More physical disks provide more spindles to service I/O. For example, in a RAID 5 database virtual disk, you can add more spindles to service I/O.

The higher the spin speed of the physical disk, the better. A 15K-RPM physical disk has one-third less rotational latency than a 10K-RPM physical disk.

NOTE The spindle count of an existing disk group can be increased using the Dynamic Capacity Expansion feature of PowerVault Modular Disk Storage Manager.

Clustering and Data Protection

Microsoft Exchange Server 2007 offers multiple options for clustering and data protection:

- Local continuous replication
- Cluster continuous replication
- Single copy clusters
- Standby continuous replication

A brief description of each option follows.

- **Local continuous replication (LCR)** – A single-server solution that uses built-in technology to create and maintain a copy of a storage group on a second set of disks that are connected to the same server as the production storage group. LCR requires a second storage array to host the passive copy. The additional storage array needs similar capacity and performance as the active copy.
- **Cluster continuous replication (CCR)** – Combines the replication and replay features in Exchange 2007 with failover features in Microsoft Windows Cluster services. CCR requires a second storage array to host the passive copy. The additional storage array needs similar capacity and performance as the active copy.
- **Single copy clusters (SCC)** – These clusters use Microsoft Windows Cluster services to manage access and node failover to a shared storage array.
- **Standby continuous replication (SCR)** – A new feature available in Exchange 2007 Service Pack 1. SCR is designed for locations that use standby recovery servers for high availability. SCR requires a second storage array to host the passive copy. The additional storage array needs similar capacity and performance as the active copy.

For more detailed information about clustering in an Exchange environment, refer to the Microsoft document, *High Availability*. Go to:

<http://technet.microsoft.com/en-us/library/bb124721.aspx>

Considering the Server Platform

The server platform consists of the server hardware and the system software.

Considering the Server Hardware Architecture

Available bandwidth depends on the server hardware. The number of buses adds to the aggregate bandwidth, but the number of HBAs sharing a single bus can throttle the bandwidth.

Calculating Aggregate Bandwidth

An important limiting factor in I/O performance is the I/O capability of the Exchange server. The aggregate bandwidth of the server to the storage array is measured in MB/s and consists of the total capability of the buses to which the storage array is connected. For example, a 64-bit PCI bus clocked at 133 MHz has a maximum bandwidth calculated by the following formula:

$$\text{PIC bus throughput (MB/s)} = \text{PCI Bus Width} \div 8 \times \text{Bus Speed}$$

$$64 \text{ bit} \div 8 \times 133 \text{ MHz} = 1062 \text{ MB/s} \approx 1 \text{ GB/s}$$

Sharing Bandwidth with Multiple HBAs

Multiple HBAs on a bus share this single source of I/O bandwidth, and each HBA might have multiple ports. As a result, the ability to drive a storage array can be throttled³ by either the server bus or by the HBAs. Therefore, whenever you configure an Exchange server or whenever you analyze I/O performance, you must know how much server bandwidth is available and which devices are sharing that bandwidth.

Host Connectivity

Dell's MD3000 and MD3000i storage arrays are designed with high availability and performance in mind. These storage arrays feature redundant power supplies, redundant RAID controller modules, redundant cache, and redundant internal architecture. To maximize this redundant technology and to prevent a single point of failure, use the following cabling guidelines in a Microsoft Exchange Server 2007 environment.

Figure 1 shows the recommended method for connecting hosts with redundant switches. Redundant switching provides the most fault tolerant configuration available for host connectivity.

³Throttle – To slow down I/O processing during low memory conditions, typically processing one sequence at a time in the order the request was received.

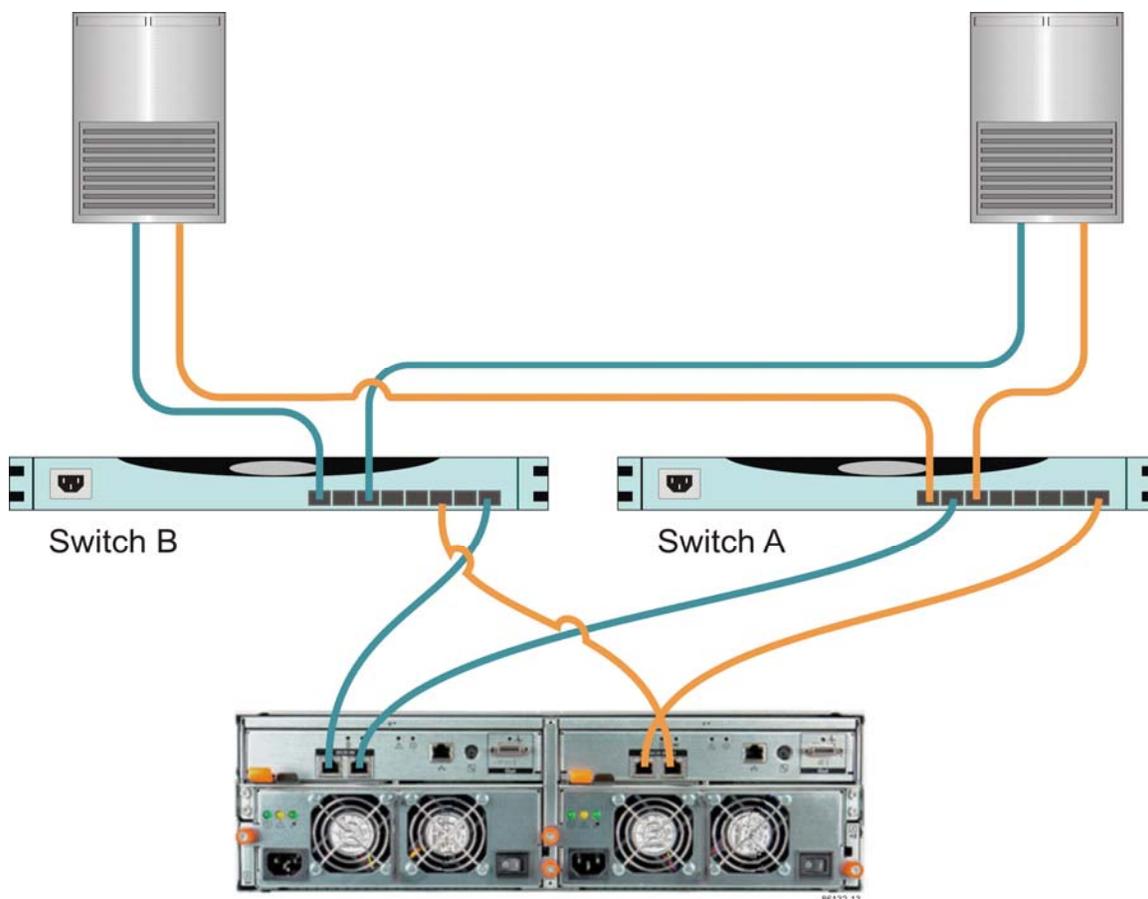


Figure 1 Redundant Switched Host Connection

Microsoft Windows Server Considerations

Aligning Host I/O with RAID Striping

On physical disks that maintain 64 sectors per track, Microsoft Windows creates the partition starting at the 64th sector. Starting at the 64th sector causes misalignment with the underlying RAID striping and allows for the possibility of a single I/O operation to span multiple segments causing performance degradation.

Use the Microsoft DiskPart.EXE utility to align the sectors of all Exchange 2007 storage virtual disks before formatting. Microsoft Exchange Server 2007 writes data in multiples of 8 KB I/O operations, and I/O operation to a database can be from 8 KB to 1 MB. Therefore, make sure that the starting offset is a multiple of 8 KB. Dell recommends starting with a value of 64 KB.

Microsoft provides the DiskPart.EXE utility as part of Windows 2003 Service Pack 1 support tools. For detailed steps about how to use DiskPart to align I/O with storage track boundaries, refer to the Microsoft document, *How to Align Exchange I/O with Storage Track Boundaries*. Go to:

<http://technet.microsoft.com/en-us/library/aag98219.aspx>

NTFS Allocation Unit Size

Configure the NTFS virtual disks that host the Exchange 2007 databases with an NTFS allocation unit size of 64 KB. The recommendation of 64 KB is based on the performance improvements seen with large sequential read operations.

Format the transaction log virtual disks with the default allocation unit size of 4 KB. Testing has shown that changing to 64 KB does not result in any increase in transaction log sequential throughput.

Setting Up the Storage Array

Factors Influencing Storage Performance

Physical disk I/O capacity is at the very heart of storage array performance. For the MD storage arrays, the number of physical disks in a configuration usually establishes the upper bound for storage array performance. Many various interacting factors determine how much of the raw performance of a group of physical disks that Microsoft Exchange Server 2007 can use. These factors include the following:

- Size of the cache
- Algorithms that manage the cache
- Number and type of host and physical disk-side channels
- How RAID parity calculations are performed
- Whether SCSI commands are queued for optimized execution by the RAID controller modules
- How RAID controller modules choose data paths

Estimating Capacity Limits

In Exchange Server 2007 environments, the number of physical disks in the physical disk group largely determines performance. The maximum IOPS from physical disk for a storage array is typically specified with a full complement of physical disks. Performance slows with fewer physical disks and can be approximated by a simple ratio.

The following formula yields an approximation only:

$$\text{Maximum I/O Rate from Maximum Number of Drives Supported} \times \text{Actual Number of Drives} \div \text{Maximum Number of Drives Supported} = \text{Maximum I/O Rate from Actual Number of Drives}$$

Calculating Optimal Segment Size

The term *segment size* refers to the amount of data written to one physical disk in a disk group before writing to the next physical disk in the disk group. Using the Modular Disk Storage Manager (MDSM) you can choose segment sizes of 128 KB or 256 KB. For example, in a RAID 5

EXCHANGE AND POWERVAULT MD3000 FAMILY

(4 + 1) disk group with a segment size of 128 KB, the first 128 KB of the LUN storage capacity is written to the first physical disk, the next 128 KB to the second physical disk, and so forth.

For a RAID 10 (2 + 2) disk group, 128 KB of an I/O would be written to each of the two data physical disks and to the mirrors. If the I/O size is larger than the number of physical disks multiplied by 128 KB, this pattern repeats until the entire I/O is completed.

The formula for optimal segment size is as follows:

$$\text{LUN segment size} = \text{LUN stripe width} \div \text{number of data physical disks}$$

For RAID 5, the number of data physical disks is equal to the number of physical disks in the disk group minus 1. For example:

$$\text{RAID 5 (4 + 1) with a 128 KB segment size} \Rightarrow (5-1) \times 128 \text{ KB} = 512 \text{ KB stripe width}$$

For RAID 1/10, the number of data physical disks is equal to the number of physical disks divided by 2. For example:

$$\text{RAID 10 (2 + 2) with a 128 KB segment size} \Rightarrow (2) \times 128 \text{ KB} = 256 \text{ KB stripe width}$$

For small I/O requests, the segment size should be large enough to minimize the number of segments (physical disks in the LUN) that must be accessed to satisfy the I/O request to minimize segment boundary crossings. For Microsoft Exchange Server 2007 virtual disks, set the segment size to 64 K or larger, so that the stripe width is at least as large as the median I/O size. Note that a segment size of 64K can only be selected using the Command Line Interface (SMcli)

Setting the Global Media Scan

The impact of Media Scan is minimal, but the extra reads do represent a finite workload. Therefore, consider the performance demands when setting Media Scan.

- When maximum performance is your objective, do not enable Media Scan.
- If performance is not a factor, media scan is valuable because it checks the drives for media errors and verifies that the volume data can be accessed. In those cases, enable Media Scan and set the scan frequency to 15 days to enable periodic scans of the surface of all LUN physical disks.

You also can enable or disable Media Scan for each virtual disk. See the following section on "Setting LUN-Specific Parameters."

Setting LUN-Specific Parameters

Setting the LUN-Specific Media Scan

One way to limit the workload caused by Media Scan, when you must use it, is to enable or disable Media Scan at the virtual disk level, rather than globally.

- You can enable Media Scan for each virtual disk.
- If the goal is to maximize the performance of a LUN, disable Media Scan for a specific virtual disk.

Setting the LUN-Specific Read Cache and Read Prefetch

Always enable read cache. Enabling read cache allows the RAID controller modules to service reads from cache for any additional read requests to the data stored within that cache.

Cache read prefetch should be disabled on all Exchange 2007 virtual disks. This prevents the RAID controller modules from reading ahead and retrieving segments of data that will most likely not be used.

Locating Disk Groups

Striping Disk Groups in RAID Controller Modules

Stripe disk groups across all available expansion enclosures to distribute I/O evenly across the drive side SAS connections. The MD3000 and MD3000i have two expansion SAS connections, one per controller. Each connection contains four SAS channels.

Figure 2 shows the striping.

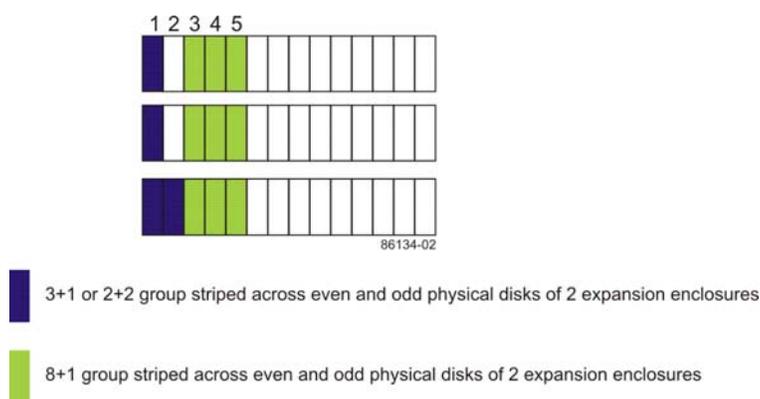


Figure 2 Striping Disk Groups in RAID Controller Modules

The blue disk group is comprised of even- and odd-numbered physical disks from each of the two available expansion enclosures. This configuration favors optimal performance: physical disks are accessed from all four channels in the SAS connection and enclosure protection is optimal.

Because the green disk group (8 + 1) has more physical disks than there are expansion enclosures, the green disk group uses multiple physical disks on all expansion enclosures. If access to an entire expansion enclosure is lost, the disk group data becomes inaccessible. However, the high component reliability and redundant design elements in the MD3000 and MD3000i storage arrays makes losing access to an entire expansion enclosure an extremely remote possibility. The choice for maximum tray protection is really targeted at human cabling errors during on-the-fly upgrades and expansion enclosure reconfigurations.

Cabling the Dell PowerVault MD3000 and MD3000i

These storage arrays can house up to 15 physical disks. Additional storage expansion can be achieved by adding up to two PowerVault MD1000 expansion units for a total of 45 drives. Figure 3 shows the maximum cascading (expansion) configuration for the

EXCHANGE AND POWERVAULT MD3000 FAMILY

duplex storage array. This cabling method is appropriate for both the MD3000 and the MD300i.

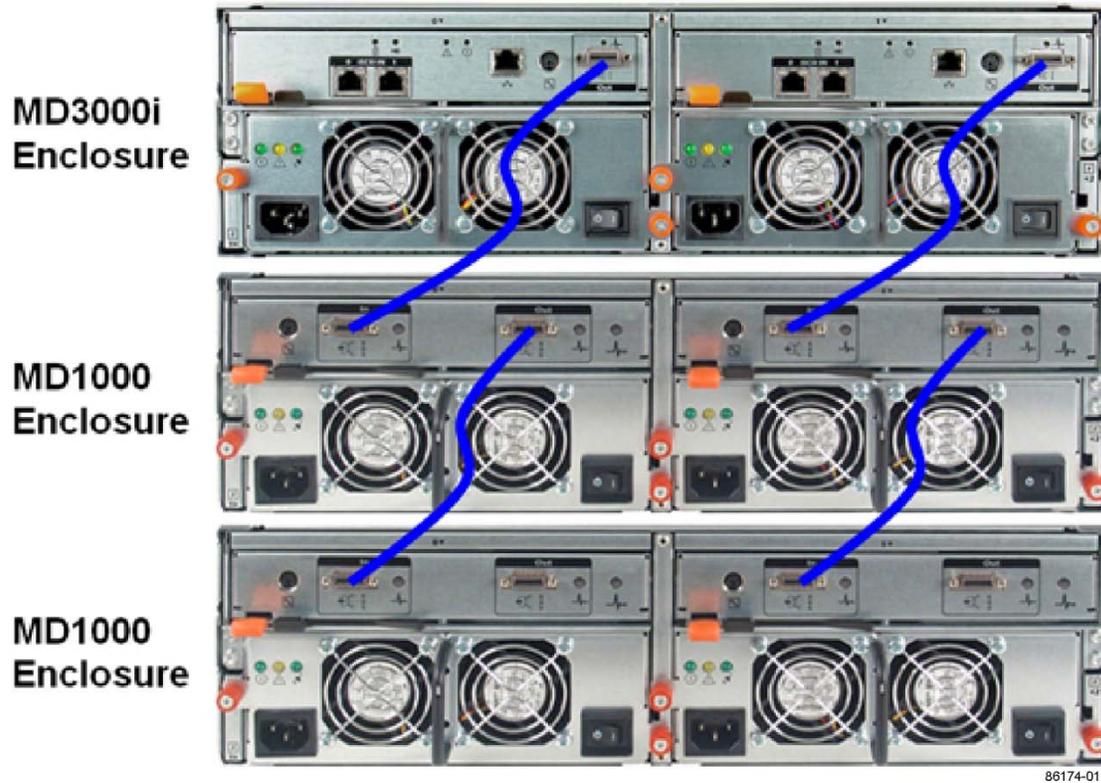


Figure 3 Cabling the MD3000i or MD3000

Using the PowerVault Modular Disk Storage Manager Performance Monitor

PowerVault Modular Disk Storage Manager provides an integrated Performance Monitor that reports the following statistics for each virtual disk in the storage array.

Table 2: Performance Monitoring Statistics⁴

Statistic	Description
Total I/Os	Since Start of this monitoring session
Read Percentage	Percent of Read I/Os
Cache Hit Percentage	Percent of Reads satisfied from cache
Current KB/sec	Since last polling interval or requested Update
Max. KB/sec	Highest value since last Start
Current I/O/sec	Since last polling interval or requested Update
Max. I/O/sec	Highest value since last Start

Use the Performance Monitor to guide the tuning process. Observe the cache hit percentage and the read/write mix for each LUN of interest while an application is running.

For detailed usage information about Performance Monitor, refer to the PowerVault Modular Disk Storage Manager online help.

⁴ These statistics are available only through the command line interface.

Summary of Performance Tuning Best Practices

Here is a summary of the recommended best practices that you should follow to obtain optimal performance from a MD3000 or MD3000i storage array in a Microsoft Exchange Server 2007 environment.

- Use all available host-side channels. Storage partitioning, switch zoning, and multi-pathing HBA drivers are all useful tools to ensure that all channels are kept active.
- Balance I/O across the dual RAID controller modules of the storage array and strive to keep both RAID controller modules busy. Locate the database virtual disk and its corresponding log virtual disk on separate paths.
- Attach cables to the expansion enclosures according to the latest *Hardware Cabling Guide*.
- Choose faster physical disks. A 15K-RPM physical disk has one-third less rotational latency than a 10K-RPM physical disk.
- Add more physical disks to the configuration for a linear increase in performance, up to the point of RAID controller module saturation. More physical disks provide more spindles to service I/O.
- For the log LUNs, use RAID 1/10 virtual disks and enable read caching and write caching with mirroring.
- For the database LUNs, use RAID 1/10 or RAID 5, and enable read caching and write caching with mirroring.
- Create disk groups across expansion enclosures to distribute I/O SAS drive connections. For details, see "Locating Disk Groups" on page 13.
- Configure the entire capacity of a disk group into a single virtual disk. Multiple LUNs on one disk group typically increases seek time penalties.
- Isolate Exchange 2007 database and the log disk I/O on separate physical disk groups. Separate the transaction log files (sequential I/O) from databases (random I/O) to maximize I/O performance and increase fault tolerance. Storing sequentially accessed files separately keeps the disk heads in position for sequential I/O, which reduces the amount of time required to locate data. From a recoverability perspective, separating a storage group's transaction logs and databases makes sure that a catastrophic failure of a particular set of physical disks does not cause a loss of both database and transaction logs.
- Choose a segment size of 64 K or larger, so that the stripe width is at least as large as the median I/O size.
- Isolate Exchange 2007 virtual disks from other physical disk-intensive applications. Sharing the storage with other applications might have a negative impact on Exchange 2007 performance.

Conclusion

This document delivers important general guidelines for fine-tuning the Windows Server, Microsoft Exchange Server 2007, and the Dell MD3000 or MD3000i storage array. To continue to improve performance, learn as much as possible about the requirements for a Microsoft Exchange Server 2007 environment.

If you have special problems that you are unable to solve, contact Dell or Dell resellers to obtain custom, on-site consulting. See the following section, “For more information on PowerVault MD3000i and MD3000 and Exchange check out:

Principled Technologies Performance Testing of PowerVault Storage and Exchange:

<http://www.dell.com/downloads/global/products/pvaul/en/exchange2007-compare-md3000i-hp.pdf>

PowerVault MD3000 ESRP Testing: <http://technet.microsoft.com/en-us/exchange/bb412164.aspx>

PowerVault MD3000i ESRP Testing:

http://www.dell.com/downloads/global/products/pvaul/en/dell_powervault_MD3000i.pdf

Contact Information.”

For more information on PowerVault MD3000i and MD3000 and Exchange check out:

Principled Technologies Performance Testing of PowerVault Storage and Exchange:

<http://www.dell.com/downloads/global/products/pvaul/en/exchange2007-compare-md3000i-hp.pdf>

PowerVault MD3000 ESRP Testing: <http://technet.microsoft.com/en-us/exchange/bb412164.aspx>

PowerVault MD3000i ESRP Testing:

http://www.dell.com/downloads/global/products/pvaul/en/dell_powervault_MD3000i.pdf

Contact Information

For more information and sales office locations, please visit the Dell web site at:

<http://www.dell.com/Storage>

References

Microsoft [2007] *High Availability* [Clustering in an Exchange environment]

<http://technet.microsoft.com/en-us/library/bb124721.aspx>

Microsoft [2006] *How to Align Exchange I/O with Storage Track Boundaries* [Microsoft's usage details on DiskPart]

<http://technet.microsoft.com/en-us/library/aag98219.aspx>

Microsoft [2007] *Mailbox Server Storage Design* [Using the Mailbox Storage Calculator for Exchange 2007 to predict storage needs]

<http://technet.microsoft.com/en-us/library/bb738147.aspx>

Microsoft [2006] *What Causes Exchange Disk I/O?* [Exchange Server 2007 I/O requirements]

<http://technet.microsoft.com/en-us/library/bb738154.aspx>

PowerVault Modular Disk Storage Manager online help [Detailed usage information about Performance Monitor]