OPTIMIZED DELL POWERVAULT STORAGE FOR MICROSOFT EXCHANGE SERVER 2007

BY SUMAN KUMAR SINGH
MAHMOUD AHMADIAN

Dell-recommended solutions for Microsoft® Exchange environments can help simplify the design and deployment of messaging infrastructures using industry-standard components. This article describes two optimized solutions using Serial Attached SCSI (SAS)–based Dell™ PowerVault™ MD1000 and PowerVault MD3000 storage.

Related Categories:

Dell PowerVault storage

Microsoft Exchange Server 2007

Serial Attached SCSI (SAS) Storage

*Visit **DELL.COM/PowerSolutions** for the complete category index.*

E-mail applications have become one of the most vital elements of enterprise IT, and the performance and availability of messaging infrastructures can be critical. Choosing the appropriate platform for messaging software such as Microsoft Exchange, however, can be a complex task. To help reduce this complexity for enterprises of all sizes, Dell has performed extensive testing and characterization of Exchange Server 2007 on standard Dell building blocks.

The Exchange Solution Reviewed Program (ESRP) was developed by Microsoft to provide a common storage testing framework that allows vendors to provide information on storage for Exchange. It integrates both the Microsoft Exchange Server Jetstress storage testing tool and solution-publishing guidelines. Dell utilizes the ESRP to test and optimize its storage for Exchange. This article presents two such recommended solutions using Serial Attached SCSI (SAS)–based Dell PowerVault MD1000 and PowerVault MD3000 storage.

Organizations should keep in mind that the ESRP is not a benchmarking program, and its tests are not designed to indicate maximum throughput for a given solution; rather, it focuses on producing recommended solutions for optimized Exchange deployments. The solutions presented in this article are only a small subset of the available ESRP solutions recommended by Dell. A complete list is available at DELL.COM/Exchange.

Dell PowerVault MD1000–based storage solution for Microsoft Exchange

The PowerVault MD1000–based storage solution for Exchange Server 2007 has been tested and validated to support 2,000 users and utilize cluster continuous replication (CCR), a high-availability clustering mechanism based on the Microsoft Cluster Service majority node set model. CCR requires an active and a passive set of Exchange Mailbox server nodes, each of which maintains its own copy of the mailbox databases. E-mail clients access the primary (active) server, and database changes to this server are sent to the secondary (passive) server in the form of log records. The log records are played on the secondary server to help keep the secondary database consistent with the primary database.

The PowerVault MD1000 is a modular disk storage expansion enclosure for Dell PowerEdge™ servers that can house up to fifteen 3.5-inch disk drives in a single 3U rackable chassis. This direct attach storage (DAS) enclosure supports both SAS and Serial ATA (SATA) disk drives. The PowerEdge Expandable RAID Controller (PERC) 5/E connects the PowerVault MD1000 storage enclosure to the server, and supports 3 Gbps SAS as the storage interconnect and PCI Express (PCIe) as the host-based interconnect.

This solution includes a primary site and a secondary site, both with a single PowerEdge 2950 server attached to a PowerVault MD1000 storage enclosure. The primary and secondary storage do not share storage array controllers or

disks. Figures 1 and 2 outline the components and configuration of this solution.

The tested user profile was for 2,000 mailboxes, assuming a 300 MB mailbox size and 0.5 I/Os per second (IOPS) for each user (0.42 IOPS per user, with room for a 20 percent increase). The storage configuration was as follows:

- **Disks 0–9:** RAID-10 volume for Exchange Information Stores
- **Disks 10–11:** RAID-1 volume for Exchange transaction logs
- **Disk 12:** Hot spare
- **Disks 13–14:** Available for future growth

This solution is designed with scalability in mind to help accommodate future growth. Organizations can easily scale it out either by daisy-chaining additional enclosures or by adding server and storage building blocks.

Servers	
Servers	One Dell PowerEdge 2950 server at each site
Processors	Two dual-core Intel® Xeon® 5160 processors at 3.00 GHz per server
Memory	16 GB of double data rate 2 (DDR2) error-correcting code (ECC) memory per server
Internal disks	Two 73 GB, 15,000 rpm Seagate ST373454SS SAS drives per server
Network interface card (NIC)	One Broadcom NetXtreme II NIC per server
RAID controller	One PERC 5/i (firmware version 1.00.01-0088) per server
Storage	
Storage	One Dell PowerVault MD1000 at each site
Disks	Fifteen 300 GB, 15,000 rpm Seagate ST33006555ss/Rev S512 SAS drives per array
RAID controller	One PERC 5/E (firmware version 1.03.10-0216) per array

Figure 1. *Components of validated Dell PowerVault MD1000–based storage solution for Microsoft Exchange*

Performance tests: I/O, backup and recovery, and log reads

In July 2007, Dell engineers conducted performance tests on the PowerVault MD1000–based storage solution for Exchange Server 2007 using the configuration described in the preceding section. Figure 3 summarizes the results. The database and transaction log I/O tests were designed to use the storage with the maximum sustainable level of Exchange I/O, to show how long the storage takes to respond to I/O under load. The data in Figure 3 is the sum of all the logical disk I/Os and the average of all the logical disks' I/O latency during the test.

The streaming backup and recovery test was designed to measure the maximum rate at which databases could be backed up. The test team measured the database and log read I/O performance metrics by running a checksum on the databases and log files. Figure 3 shows the average rate for a single database file.

The log read test was designed to measure the maximum rate at which the log files could be played against the databases. Figure 3 shows the average rate for 500 log files played in a single storage group, where each log file was 1 MB in size.

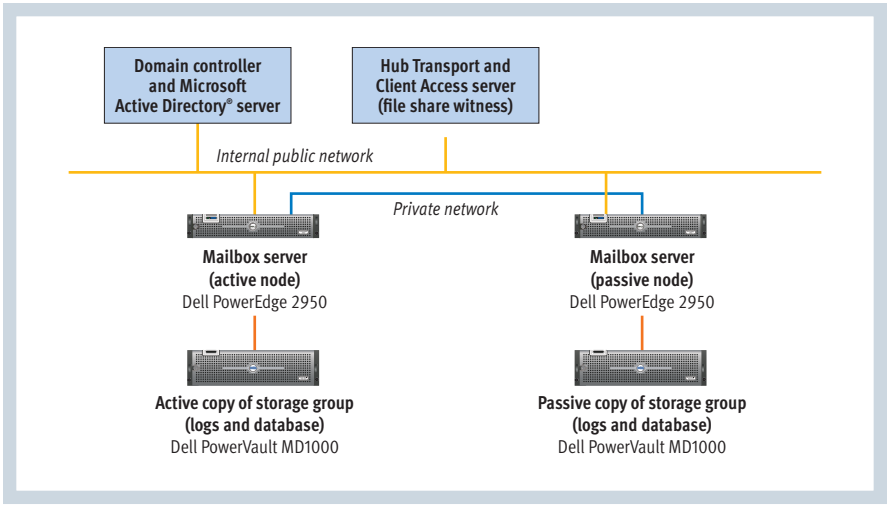

Figure 2. *Configuration of validated Dell PowerVault MD1000–based storage solution for Microsoft Exchange*

Database I/O	
Average database disk transfers/sec	1,435
Average database disk reads/sec	745
Average database disk writes/sec	690
Average database disk read latency	16 ms
Average database disk write latency	4 ms
Transaction log I/O	
Average log disk writes/sec	477
Average log disk write latency	0 ms
Streaming backup and recovery	
Average 1 MB reads/sec per storage group	57
Total 1 MB reads/sec	286
Log reads	
Average time to play one log file	0.36 sec

Figure 3. *Performance results for the validated Dell PowerVault MD1000–based storage solution for Microsoft Exchange*

Servers	
Servers	Two Dell PowerEdge 2950 servers
Processors	Two dual-core Intel Xeon 5160 processors at 3.00 GHz per server
Memory	16 GB of DDR2 ECC memory per server
Internal disks	Two 73 GB, 15,000 rpm Seagate ST373454SS SAS drives per server
NIC	One Broadcom NetXtreme II NIC per server
RAID controller	One PERC 5/i (firmware version 1.00.01-0088) per server
Storage	
Storage	Dell PowerVault MD3000 (firmware version 06.17.77.60), with two PowerVault MD1000 expansion enclosures
Disks	Fifteen 300 GB, 15,000 rpm Seagate ST33006555ss/Rev S512 SAS drives per array
RAID Controller	One PERC 5/E (firmware version 00.10.49.00) for the PowerVault MD3000

Figure 4. *Components of validated Dell PowerVault MD3000–based storage solution for Microsoft Exchange*

Dell PowerVault MD3000–based storage solution for Microsoft Exchange

The PowerVault MD3000–based storage solution for Exchange Server 2007 has been tested and validated to support 3,000 users and utilize single copy clustering (SCC), a high-availability clustering mechanism based on Microsoft Cluster Service shared storage. SCC requires an active and a passive set of Exchange Mailbox server nodes sharing the mailbox databases. During normal operations, e-mail clients access the primary (active) server. If the primary node fails, the secondary (passive) node takes over as the active node. Compared with previous versions of Exchange software, Exchange Server 2007 has enhanced deployment setup and management functionality for this clustering model.

The PowerVault MD3000 is a modular disk storage system that can expand with up to two PowerVault MD1000 storage enclosures. The PERC 5/E connects the PowerVault MD3000 to the server, and supports 3 Gbps SAS as the storage interconnect and PCIe as the host-based interconnect.

This solution includes two PowerEdge 2950 servers attached to a PowerVault MD3000 storage system. The primary (active) and secondary (passive) nodes use the same configuration, with the primary node providing client access to the Exchange Information Store. The storage is shared between these two nodes. Figures 4 and 5 outline the components and configuration of this solution.

The tested user profile was for 3,000 mailboxes, assuming a 1 GB mailbox size and 0.5 IOPS per user (0.42 IOPS per user, with room for a 20 percent increase). The storage configuration in each of the three storage enclosures was as follows:

- **Disks 0–9:** RAID-5 volume for Exchange Information Stores
- **Disks 10–11:** RAID-1 volume for Exchange transaction logs
- **Disk 12:** Hot spare
- **Disks 13–14:** Available for future growth

Like the PowerVault MD1000–based solution, organizations can scale this solution out by adding server and storage building blocks.

Performance tests: I/O, backup and recovery, log reads, and RAID level

In August 2007, Dell engineers conducted performance tests on the PowerVault MD3000–based storage solution for Exchange Server 2007 using the configuration described in the preceding section. Figure 6 summarizes the results of these tests, which had the same design and goals as those performed on the PowerVault MD1000–based solution.

The ESRP focuses on storage testing to help address performance and reliability issues with storage design; however, storage is not the only factor to take into consideration when designing a scalable Exchange solution. For example, the PowerVault MD3000–based storage solution for Exchange uses RAID-5 containers for the Exchange database, which can be effective for environments with very large mailboxes and/or

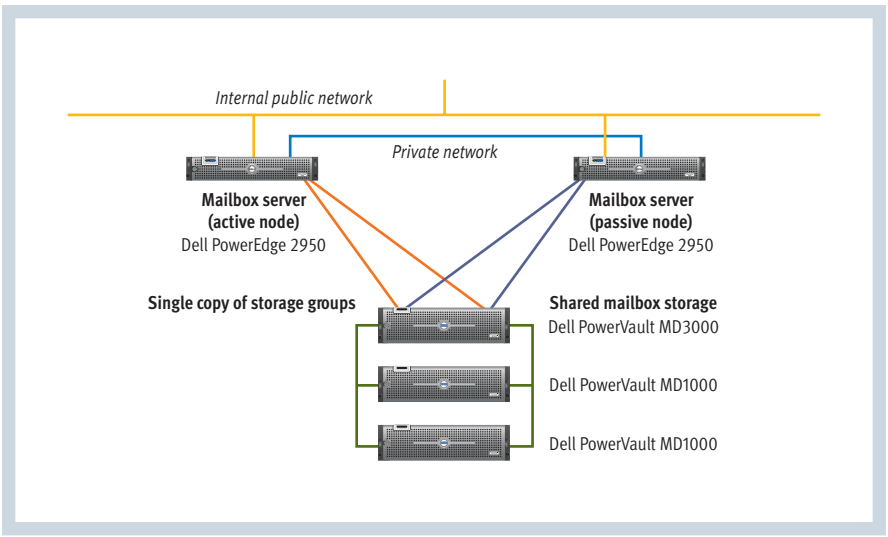

Figure 5. *Configuration of validated Dell PowerVault MD3000–based storage solution for Microsoft Exchange*

"E-mail applications have become one of the most vital elements of enterprise IT, and the performance and availability of messaging infrastructures can be critical."

low IOPS requirements. RAID-10 typically offers the highest performance with high levels of protection, but only leaves half of the RAID group's capacity usable. RAID-5 offers a higher usable capacity per RAID group than RAID-10, but lower I/O throughput.

There are two reasons for this reduced throughput. First, under normal operating conditions, RAID-5 requires the extra work of generating and updating parity information, which is what makes RAID-5 fault tolerant and RAID recovery possible. Second, following a disk failure, RAID-5 containers must restore data and calculate its parity in addition to serving normal I/O requests, which can significantly reduce the performance during rebuild operations. In addition, a RAID-5 volume can only tolerate a single disk failure, while RAID-10 can survive up to *n* disk failures for an *n* + *n* RAID-10 volume.

As shown in Figure 7, performance can suffer significantly during a RAID-5 rebuild cycle. In an optimal state—before failure—the RAID-5 database disks provided 1,709 IOPS with a read latency of 11 milliseconds and a write latency of 9 milliseconds. In a degraded state, the I/O throughput was reduced to 1,545 IOPS and the database read latency increased to 14 milliseconds. During the rebuild cycle, the throughput was further reduced to 1,116 IOPS—approximately 35 percent less than the throughput in the optimal state—with a read latency of 18 milliseconds and a write latency of 16 milliseconds.

A variation of the PowerVault MD3000–based solution would use RAID-10 for the Exchange databases, which could provide higher throughput than RAID-5. However, this approach would also reduce the total size of the database logical units from 7,344 GB to approximately 4,080 GB, thereby limiting the user mailbox size to 650 MB rather than 1 GB.

Optimized storage for Microsoft Exchange

SAS-based Dell PowerVault systems provide a standards-based storage platform for Microsoft Exchange deployments. By taking advantage of the Dell-recommended PowerVault MD1000– and PowerVault MD3000–based solutions described in this article, enterprises can help simplify the design and deployment of their Exchange Server 2007 infrastructure while helping optimize messaging performance.

Suman Kumar Singh *is a lead engineer on the Dell End-to-End Solutions team. He specializes in messaging systems architecture and sizing, and led the release of the Dell Exchange 2003 Advisor tool. His other interests include storage area networks (SANs), virtualization, and security, and he has published and presented several papers at industry conferences. Suman has a master's degree from the University of Texas at El Paso.*

Mahmoud Ahmadian *is an engineering consultant on the Dell End-to-End Solutions team. His technical interests include clustering technologies, databases, and messaging systems performance. He has an M.S. in Computer Science from the University of Houston, Clear Lake.*

	Optimal	Degraded	Rebuilding
Database IOPS	1,709	1,545	1,116
Read latency	11 ms	14 ms	18 ms
Write latency	9 ms	9 ms	16 ms

Figure 7. *Test results for RAID-5 performance in the validated Dell PowerVault MD3000–based storage solution for Microsoft Exchange*

Database I/O	
Average database disk transfers/sec	1,709
Average database disk reads/sec	887
Average database disk writes/sec	822
Average database disk read latency	12 ms
Average database disk write latency	11 ms
Transaction log I/O	
Average log disk writes/sec	518
Average log disk write latency	2 ms
Streaming backup and recovery	
Average 1 MB reads/sec per storage group	22.4
Total 1 MB reads/sec	336
Log reads	
Average time to play one log file	0.50 sec

Figure 6. *Performance results for the validated Dell PowerVault MD3000–based storage solution for Microsoft Exchange*

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Microsoft Exchange Server 2007 solutions on Dell:
DELL.COM/Exchange