

Storage Consolidation with the Dell PowerVault MD3000i iSCSI Storage

By Dave Jaffe
Dell Enterprise Technology Center
and
Kendra Matthews
Dell Storage Marketing Group



Dell | Enterprise Technology Center
delltechcenter.com

September 2007

Contents

Executive Summary	3
Introduction	4
The Dell™ PowerVault™ MD3000i	5
The Servers and their Workloads	9
Exchange 2007	9
SQL Server 2005	9
Web Serving	9
File Serving	10
Backup	10
The Storage Consolidation Simulation	11
Conclusions	14

Tables

Table 1: Dell PowerVault MD3000i Features	6
Table 2: PowerVault MD3000i Disk Layout.....	10
Table 3: Storage Consolidation Test Results. Test performed 9/2007.	13

Figures

Figure 1: Dell PowerVault MD3000i Front View	5
Figure 2: Dell PowerVault MD3000i Rear View.....	5
Figure 3: Modular Disk Storage Manager Summary Screen	6
Figure 4: Modular Disk Storage Manager iSCSI Tab	7
Figure 5: Modular Disk Storage Manager Host Topology Screen.....	8
Figure 6: The Storage Consolidation Test Setup (MD3000i Storage Array shown with bezels removed)	11
Figure 7: Symantec Backup Exec Screen Shot	12

For Further Information and Discussion

Visit our Dell TechCenter wiki at <http://delltechcenter.com/page/PowerVault+MD3000i+Storage+Consolidation+Study> to discuss this paper, and to contribute your experiences in using the Dell PowerVault MD3000i iSCSI storage for your business.

Talk Back

Tell us how the Dell Scalable Enterprise Technology Center can help your organization better simplify, utilize, and scale enterprise solutions and platforms. Send your feedback and ideas to Enterprise_Techcenter@dell.com or visit us at dell.com/TechCenter.

Executive Summary

To demonstrate the effectiveness of the PowerVault MD3000i in consolidating storage from multiple servers, a test was run in which the application data from servers hosting an Exchange 2007 cluster, a web server, a database server and a file server was stored on a single PowerVault MD3000i with 45 drives. The applications were driven for 8 hours, with two incremental Exchange backups performed as well. The PowerVault MD3000i was able to handle the I/O from all these servers, supporting 500 users on a fully redundant Exchange cluster, web and database servers hosting online e-commerce at a 2 billion orders per year clip, and a file server delivering over 700 files. In total the MD3000i supplied 387 GB during the 8-hour test.

Introduction

Customers continue to look for cost effective ways to streamline their IT infrastructure and optimize data management. Storage consolidation through the deployment of a storage area network (SAN) improves utilization and simplifies management. In the past, to get the benefits of wide scale storage consolidation customers had to invest in a Fibre Channel infrastructure. Today there is an alternative, Internet SCSI (iSCSI), which transmits data packets over standard Ethernet networking. With iSCSI technology, customers can leverage existing networking expertise and equipment to simplify their implementation of a storage network, removing the traditional barriers to storage consolidation.

The Dell PowerVault MD3000i is an iSCSI SAN that can consolidate up to sixteen hosts, expands to support up to 18TB of data and provides wizard based installation, intuitive management, and advanced data protection software.

The PowerVault MD3000i system's modular expandability provides the flexibility to add capacity as needed. The storage array can house up to fifteen 3.5-inch serially-attached SCSI (SAS) disk drives. Its capacity can be easily expanded by adding up to two PowerVault MD1000 expansion enclosures, for a total of 45 disk drives.

The PowerVault MD3000i delivers a seamless suite of intuitive, intelligent storage-management software capabilities. The Modular Disk Storage Manager automatically configures the system for optimal performance and availability. The Recovery Guru tool diagnoses system problems and helps to determine an appropriate recovery procedure. Optional snapshot and virtual disk copy features that support real-time backups of the entire data environment are designed to enhance data protection.

For high availability, the PowerVault MD3000i supports redundant active/active controllers, management ports and power/cooling to provide real security at the hardware level and can automatically rebuild a failed drive using a global hot-spare drive.

To demonstrate storage consolidation with the MD3000i, the application data of five Dell PowerEdge servers was placed onto a single MD3000i. The applications running on the five servers – the active and passive nodes of an Exchange cluster continuous replication configuration, a web server, a database server, and a file server – were then stressed using 4 separate workload driver programs for an 8-hour period. Additionally, backups were run against the passive Exchange node twice during the test.

The next section describes the PowerVault MD3000i that was used in the test, Section 4 describes the Exchange, web serving, SQL Server, file serving and backup workloads used in the simulation including details of how the data was organized on the disks, and Section 5 shows the results of the consolidation test.

The Dell™ PowerVault™ MD3000i

The PowerVault MD3000i Modular Disk Storage Array (shown in Figure 1) is a SAS-based disk array with one or two embedded iSCSI controllers. The redundant controllers are seen in the PowerVault MD3000i rear view (Figure 2). Each controller includes two 1 Gb/s iSCSI ports, as well as an Ethernet management port and a SAS expansion port.



Figure 1: Dell PowerVault MD3000i Front View

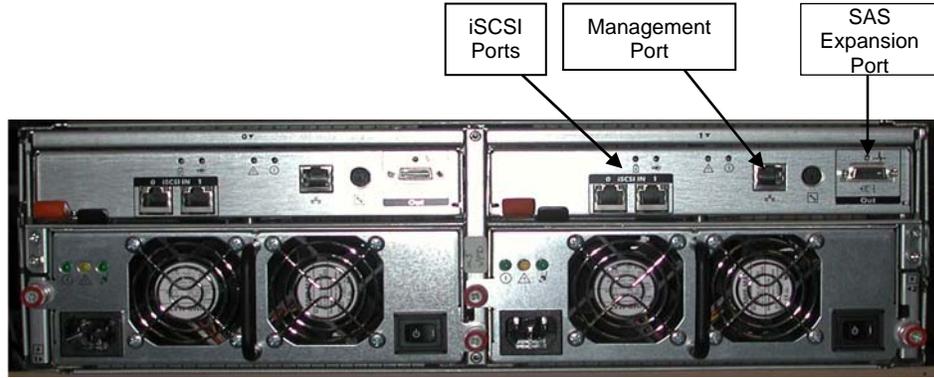


Figure 2: Dell PowerVault MD3000i Rear View

In addition to the integrated disk array the PowerVault MD3000i can be expanded with up to two PowerVault MD1000 Disk Expansion Enclosures for a total of 45 drives. For this test a PowerVault MD3000i configuration with two controllers and three enclosures containing 45 146GB 15K RPM SAS drives was used. PowerVault MD3000i features are summarized in Table 1

Disks Available	15K SAS drives available in 36GB ¹ , 73GB, 146GB or 300GB 10K SAS drives available in 146GB, 300GB or 400GB
Size including 1 enclosure with 15 disks(1.75" Rack Units)	3U
RAID Cache Size	512 MB per controller
Maximum Number of Disks (including 2 expansion enclosures)	45
List Price as tested (as of 9/17/07; includes 3-year Gold Enterprise Support)	\$45,908

¹For hard drives, GB means 1 billion bytes and TB equals 1 trillion bytes; actual capacity varies with preloaded material and operating environment and will be less.

Table 1: Dell PowerVault MD3000i Features

The PowerVault MD3000i is managed through the Dell Modular Disk Storage Manager (MDSM) which can be installed on any server that is networked to the MD3000i management port. Through this tool virtual disks (also known as storage logical units or LUNs) are created and made accessible to specified hosts. The Summary tab (Figure 3) provides an overview of hosts and virtual

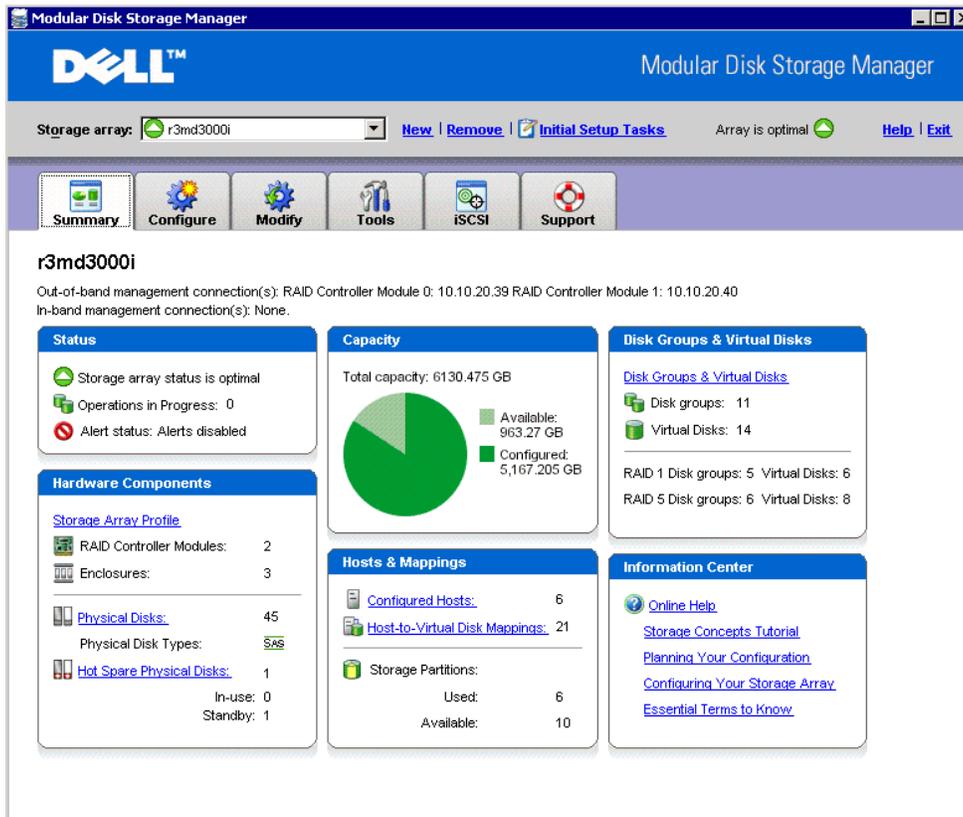


Figure 3: Modular Disk Storage Manager Summary Screen

disks as well as current storage array status. The array name and IP address of the management ports may be set under the Tools tab. Management of the MD3000i iSCSI interface including the IP addresses of the 4 iSCSI ports, authentication and target discovery is handled under the iSCSI Tab (Figure 4).

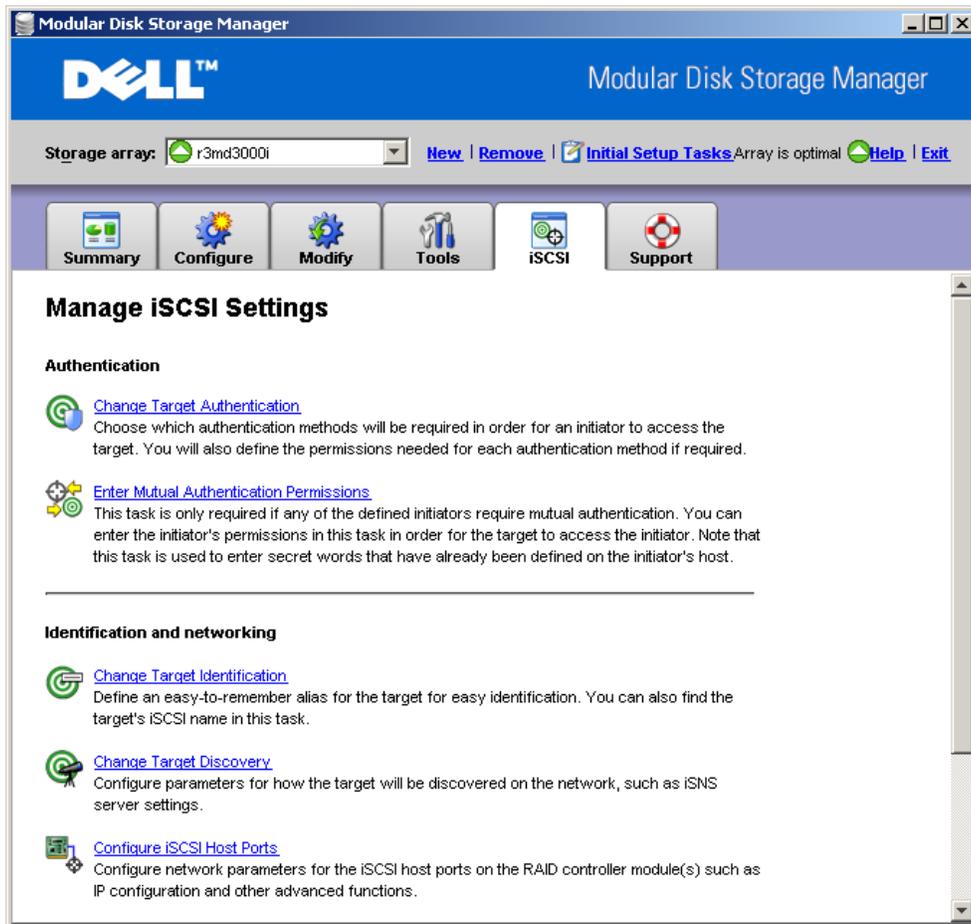


Figure 4: Modular Disk Storage Manager iSCSI Tab

The Configure Host Access (Manual) function under the Configure tab can be used to give a host server access to LUNs created on the MD3000i. The Edit Host Topology function under the Modify tab is used to add the iSCSI Initiator for each host.

The LUNs that will hold the servers' application data may be created before or after the servers are added to the Host Topology. Once a LUN is created the iSCSI initiator for the host operating system (Microsoft iSCSI Initiator on all servers in this test) is used to configure a multipathed connection to the LUN for load balancing and failover. Complete details of the multipath I/O (MPIO) creation are shown at <http://delltechcenter.com/page/MPIO+to+PowerVault+MD3000i+with+Microsoft+iSCSI+Initiator>.

The servers given access to the MD3000i in this test are shown in Figure 5.

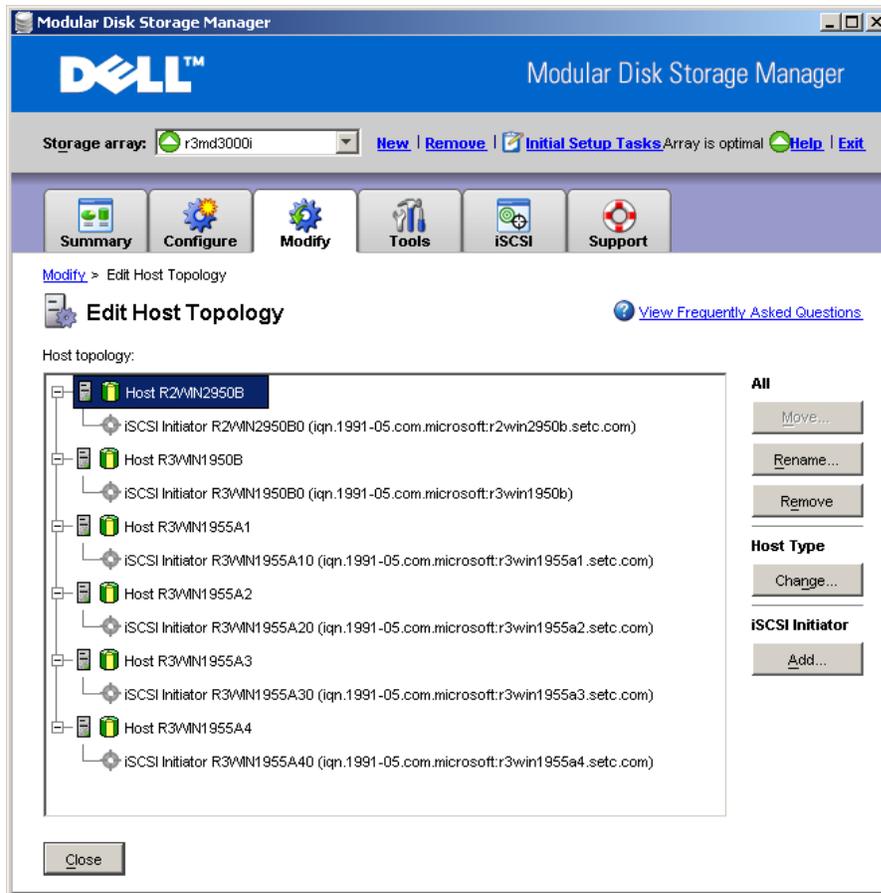


Figure 5: Modular Disk Storage Manager Host Topology Screen

The Servers and their Workloads

In this section the five separate servers whose data was consolidated on the PowerVault MD3000i will be described, along with the Exchange backup used.

Exchange 2007

Exchange Server 2007 cluster continuous replication (CCR) was installed on two PowerEdge 1955 server blades, with one blade serving as the active node and one as the passive node. Exchange 2007 CCR works by shipping transaction logs from the active node to the passive node, thus keeping the two nodes synchronized. In the event of a loss of the active node the Exchange users are automatically failed over to the passive node. Two 4-disk RAID 5 Data LUNs and a 2-disk RAID 1 Log LUN were created for the active node on Enclosure 0 and for the passive node on Enclosure 1 as shown in Table 2. (Full fault tolerance would require the passive node LUNs to be on a separate storage array altogether).

Microsoft Load Generator 2007 was used to simulate 500 heavy email users running against the Exchange 2007 cluster. Load Generator 2007 is a tool provided by Microsoft to simulate email users and is typically run over an eight hour period to simulate a complete workday. Load Generator comes with several default profiles and for this test the Heavy Exchange 2007 online user profile was selected which means that each user averages 46 actions per day. Load Generator reports results in terms of average latency for each of the different types of email operations.

SQL Server 2005

The SQL Server workload represented the backend of an online e-commerce application simulated by the Dell DVD Store database test, available at linux.dell.com/dvdstore. SQL Server 2005 was installed on a third PowerEdge 1955 blade. The Large version of the DVD Store database, over 100 GB of data, was built on two 4-disk RAID 5 Data LUNs and a 2-disk RAID 1 Log LUN. The SQL Server driver program, included with the DVD Store kit, was used to simulate a constant load of users logging in to the store, searching for DVDs by actor, title or category, and purchasing DVDs, for the 8 hour duration of the test.

Web Serving

To provide a web serving application for this test, an ASP.NET-based web application from the Dell DVD Store kit running under Microsoft Internet Information Server (IIS) was used in front of a second SQL Server database, with both IIS and SQL Server running on a 4th PowerEdge 1955 blade server. The web application provides a full set of login, search and purchase screens for the online DVD Store. A separate web driver program from the DVD Store kit was used to simulate multiple users searching for and ordering DVDs from their web browsers. This workload was also run for the duration of the test.

The backend SQL Server database used for the web serving application also used the Large version of the DVD Store. The data and log LUNs (labeled "SQLW" in Table 2) were of the same type as the SQL Server test (two 4-disk RAID 5 Data LUNs and a 2-disk RAID 1 Log LUN) and were created using the same disk group as the SQL Server test. In other words, the same set of 10 disks was used simultaneously by two 100GB databases for an 8-hour run.

File Serving

For the file serving workload a 2-disk RAID 1 LUN was created and assigned to a PowerEdge 1950 running Windows 2003 Server. The LUN was assigned a disk letter and shared as a Windows 2003 Server file share. A set of 41 files, with sizes varying from 1 MB to 3 GB, was copied to it. This file share was then mounted to another PowerEdge server using the standard Window Explorer map network drive function. To simulate employee use of the file share, the files were copied back and forth from this server to the file share using a simple Windows command shell script. Two scripts were run in parallel. Delays were placed in the script so that each script copied roughly 40 files per hour.

Backup

Symantec Backup Exec™ 11d for Windows Server was used to back up the Exchange database (see "Protecting Critical Enterprise Applications with Symantec Backup Exec 11d" by Charles Butler in Dell Power Solutions, May 2007, <http://www.dell.com/downloads/global/power/ps2q07-20070408-Symantec.pdf>). Backup Exec enables fast, comprehensive backup and recovery to disk and/or tape. After a full backup of the Exchange database is performed, incremental backups are taken from the passive node of the Exchange cluster so as not to impact the active Exchange users.

The complete disk layout used in the test is shown in Table 2. In addition to the application data LUNs described above, Disk 0 in Enclosure 0 was assigned as a hot spare for the entire array and a ten-disk RAID10 snapshot repository was created on Enclosures 1 and 2 to provide snapshot protection from both disk and enclosure failure.

	Disk 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
0	Hot Spare	Avail	Files RAID 1	Ex Act Logs RAID 1	Exchange Act Data 1 RAID 5	Exchange Act Data 2 RAID 5									
Enclosure 1	Snapshots RAID 10				Ex Pas Logs RAID 1	Exchange Pas Data 1 RAID 5	Exchange Pas Data 2 RAID 5								
2	Snapshots RAID 10				SQLW Logs RAID 1	SQLW Data 1 RAID 5	SQLW Data 2 RAID 5								
					SQL Logs RAID 1	SQL Data 1 RAID 5	SQL Data 2 RAID 5								

Table 2: PowerVault MD3000i Disk Layout

The Storage Consolidation Simulation

The lab setup used in the 8-hour simulation is shown in Figure 6. The servers, including the four PowerEdge 1955 blades running Exchange active and passive nodes, web serving and SQL Server, and the PowerEdge 1950 acting as the file server, were connected to the PowerVault MD3000i array by two subnets, 10.10.20 and 10.10.22, using the Microsoft iSCSI Software Initiator to provide multiple connection paths to the array for load balancing and failover. For simplicity the PowerEdge 1950 also hosted the MDSM management console for the PowerVault MD3000i. (In general the management server will be on a separate network from the dedicated iSCSI networks). The server running Backup Exec doesn't need to be connected directly to the MD3000i – it pulls the incremental backups from the iSCSI storage through the Exchange passive node.

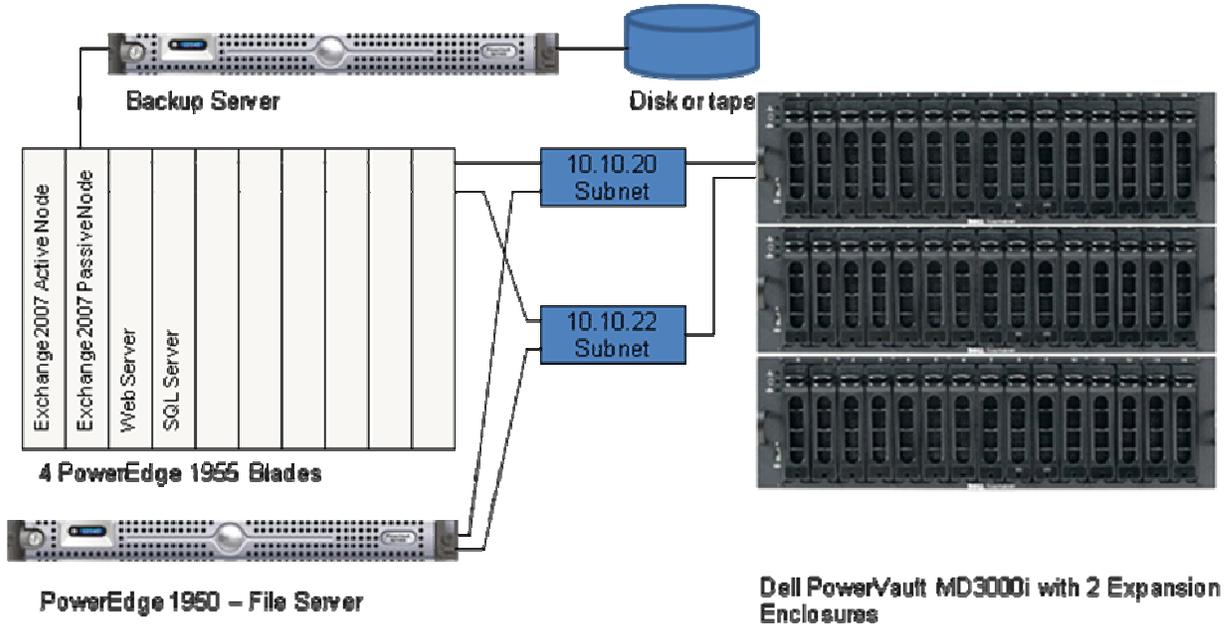


Figure 6: The Storage Consolidation Test Setup (MD3000i Storage Array shown with bezels removed)

The Exchange, SQL Server, web serving and file serving workloads were started against the five servers as described in Section 4 and run for 8 hours to represent a typical 8 am – 4 pm work day. Each workload driver reported application-specific metrics (Exchange tasks, orders per minute, etc.). Windows Performance Monitor was used to record average throughput (in bytes per second) and other disk parameters. Incremental backups were performed with Backup Exec at times during the run corresponding to 11:30 am and 3:30 pm during the simulated day (the backups dated “9/12/2007” in Figure 7).

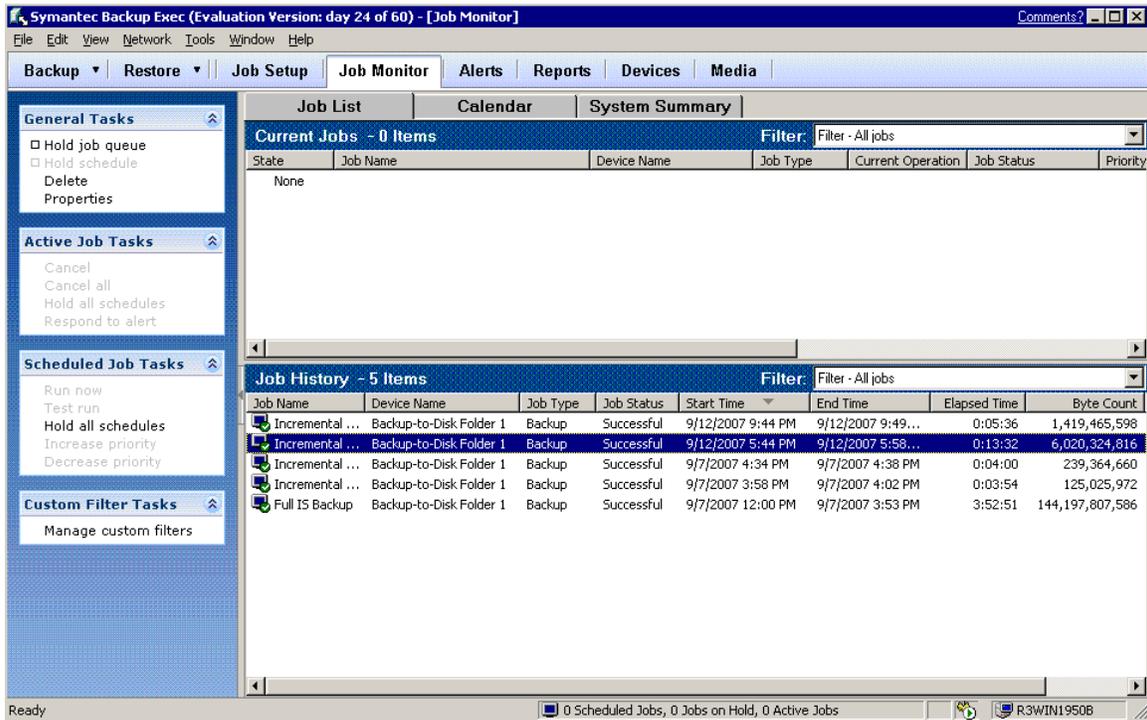


Figure 7: Symantec Backup Exec Screen Shot

The performance of the 5 servers over the 8-hour simulation is tallied in Table 3. The 500 heavy profile Exchange 2007 users completed 23,832 tasks (reads, replies, calendar updates, etc.) on the active Exchange node. The Exchange transaction log files were shipped continuously to the passive Exchange node, keeping the two nodes in sync. The passive node also handled the two backups described above. The web server handled 0.77 million orders during the 8-hour day while the SQL Server database handled 1.1 million orders (together this is a 24/7 run rate of over 2 billion orders per year). The file share copied 705 files back and forth for a total transfer of 54 GB. In all the MD3000i supplied 387 GB of data during the 8-hour run.

Server/Workload	Total Work in 8-hour Simulation	Total Transferred (GB)
r3win1955a1 Exchange 2007 active node	500 users, 23832 tasks	42
r3win1955a2 Exchange 2007 passive node	500 users, 23832 tasks, plus 2 backups	82
r3win1955a3 Web Serving	0.77 million orders	61
r3win 1955a4 SQL Server 2005	1.1 million orders	148
r3win1950b CIFS	705 files copied	54
Total	-	387

Table 3: Storage Consolidation Test Results. Test performed 9/2007.

Conclusions

The PowerVault MD3000i is a cost effective, simple to manage iSCSI storage array that can be used to consolidate the storage of several servers, providing higher storage utilization, integrated snapshots and a single management interface. To demonstrate this, the application storage requirements of five servers were built on a PowerVault MD3000i with 45 disks. The five applications running on the servers were then exercised with load drivers for an 8-hour period. Additionally, two incremental backups were made against the passive Exchange node during the run. The MD3000i was able to handle the I/O from all these servers, supporting 500 users on a fully redundant Exchange cluster, web and database servers hosting online e-commerce at a 2 billion orders per year clip, and over 700 files copied. In total the MD3000i supplied 387 GB during the 8-hour test.

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, PowerEdge and PowerVault are trademarks of Dell Inc. Exchange and SQL Server are registered trademarks of Microsoft Inc. Intel and Xeon are registered trademark of Intel Corp. Red Hat is a registered trademark of Red Hat Inc. Linux is a registered trademark of Linus Torvalds. 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

Pricing, specifications, availability, and terms of offers may change without notice. Taxes, fees, shipping, handling and any applicable restocking charges are extra, and vary.

©Copyright 2007 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Information in this document is subject to change without notice.