

LAB VALIDATION REPORT

Quantum DXi7500 Deduplicated backup for the Enterprise

By Tony Palmer
With Brian Garrett and Lauren Whitehouse

December, 2008

Table of Contents

Table of Contents	i
Introduction	1
<i>Background.....</i>	<i>1</i>
<i>Quantum Disk-based Backup</i>	<i>2</i>
ESG Lab Validation.....	5
<i>Policy-based deduplication operation.....</i>	<i>5</i>
<i>Policy-based deduplication performance.....</i>	<i>8</i>
<i>Remote Replication</i>	<i>11</i>
<i>Path-to-tape</i>	<i>13</i>
<i>Manageability</i>	<i>17</i>
ESG Lab Validation Highlights.....	21
Issues to Consider	21
ESG Lab's View	22
Appendix.....	23

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by Quantum.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

Introduction


Quantum's portfolio of disk-based backup appliances has been improving performance and reliability for customers since 2002. In 2007, ESG Lab tested Quantum's DXi Series of disk-based backup and replication appliances designed for mid-sized environments and remote offices. This report examines in detail Quantum's new enterprise-focused DXi7500 platform, featuring "policy-based deduplication" which allows users to choose between deduplication methods. ESG lab also explores the DXi7500's role in a full edge-to-core data protection solution, focusing on bandwidth-optimized remote replication, physical tape integration as well as performance and scalability optimized for data requirements.

Background

A growing number of organizations are turning to disk-based backup and recovery methods to improve the speed and reliability of backup and restore operations for ever-expanding volumes of data to be protected. Disk-based systems are inherently faster and more reliable than tape, but they can be significantly more expensive. Data deduplication can effectively bring down the expense of disk-based storage and make it more practical. Block-based data deduplication reduces capacity requirements by examining data at a sub-file level and ensuring that only unique data is written to disk. In other words, when data that has already been seen before is ingested, the appliance stores a pointer to the original data rather than storing a new copy of the data. Data deduplication eliminates the need to store multiple copies of the same data over time, dramatically reducing backup to disk capacity requirements.

ESG research indicates that no single concern dominates when users are considering data deduplication technology. As shown in Figure 1, compatibility with existing solutions, ease of integration and management, performance, and scalability all have significant mindshare with users in the enterprise. In order to succeed at the enterprise level, data deduplication appliances must provide cost-effective scalability, performance and flexibility.

FIGURE 1. DATA DEDUPLICATION EVALUATION CRITERIA



While deduplication can reduce the cost of the raw storage required to store backups on disk by enabling organizations to retain backup data for longer periods of time, integration within the organization's ecosystem is crucial—robust management, tape integration, flexible performance and deduplication options are all important considerations for enterprises. The DXi7500 is being positioned by Quantum as an enterprise-class answer to these concerns.

Quantum Disk-based Backup

ESG Lab has had the unique opportunity to follow the continued evolution of Quantum's disk-based backup target systems since first testing the DX100 in 2004. In December 2006, Quantum introduced its DXi3500 and DXi5500 midrange disk backup appliances with deduplication and replication; ESG Lab tested these just over a year ago. In last year's report, ESG Lab examined overall performance, deduplication rates, and the effectiveness of WAN-based remote replication, which leverages deduplication to reduce bandwidth requirements. ESG Lab found that the Quantum approach to this technology, which combines variable-length deduplication and built-in hardware compression, delivers dramatic savings in disk capacity when storing backup data on disk and provides similar reductions in network bandwidth requirements when replicating deduplicated data over a WAN connection.

FIGURE 2. EFFECT OF DATA DEDUPLICATION OVER TIME




Figure 2 illustrates the effects of data deduplication over time as observed by ESG Lab during testing in 2007. Multiple full backups were run and the projected data reduction over thirty days was calculated. The capacity savings achieved by the DXi5500 is represented by the percentage of the total volume of protected data sent to the system by the backup application actually stored to disk. In this example, the amount of data stored on disk after 30 simulated days was 115 GB—about 4% of the total 2,699 GB of data backed up after thirty days. This represents a 96% reduction in data storage requirements. Also confirmed in ESG Lab's 2007 evaluation of the

DXi Series appliances were the products' seamless integration with existing backup environments. Not only was the solution easy to implement, but its easy to use management console made the configuration process very straightforward. The flexibility of the DXi Series was also noteworthy, with support for many different backup applications, presentations (NAS and VTL, iSCSI and Fibre Channel), and a wide variety of tape devices.

Since organizations implement disk in the backup process to accelerate backups and restores, performance was also tested. Quantum didn't disappoint, posting impressive backup throughput and near-instant recovery of files. Organizations looking to reduce their reliance on physical tape are also interested in strategies for disaster recovery once a backup-to-disk approach is deployed. The DXi Series was put through a series of tests for VTL-to-VTL replication over a WAN connection. Leveraging the deduplication system brought significant optimization to WAN bandwidth, making it possible to transport large amounts of data over distance quickly and efficiently.

ESG Lab's validation of the DXi Series in 2007 confirmed Quantum's claims of simplicity, flexibility, efficiency, and high performance. The highlighted features made a strong argument for backup-to-disk and demonstrated how Quantum's deduplication and replication solutions can provide higher levels of service with maximum cost efficiency.

The Quantum DXi7500

The DXi7500 is a disk backup solution for distributed and enterprise environments designed for larger midrange sites and corporate data centers. It uses the same variable-length block deduplication approach and the same replication technology as Quantum's other deduplication appliances. Raw capacity scales from 12 TB (9 TB usable) to 240 TB (180 TB usable, with usable capacity expanding soon to 220TB), allowing users to grow their system in manageable increments while preserving data, and it offers ingest performance of up to 4 TB/hour. Quantum has also incorporated policy-based deduplication in the DXi7500, a unique feature of the company's enterprise deduplication technology that allows users to choose different modes of deduplication for different data sets.

Background

There are two basic approaches to carrying out deduplication: de-duplicating data during the ingest process and waiting until after the backup. Users care about which approach is used because all deduplication includes some level of overhead. If deduplication occurs during ingest, that overhead occurs during the backup window and may have an impact on performance of the backup. If deduplication is delayed until after the backup is finished, the overhead is incurred outside the backup window and performance of the backup is unaffected. Until the introduction of the DXi7500, all vendors offered either one approach or the other—Quantum's DXi7500 is the first product to provide users with a choice, letting them decide when to perform deduplication for different data sets or backup jobs as well as giving them the ability to turn it off altogether. A brief summary of each mode and the associated use cases follows.

Deduplication during ingest: Quantum calls this policy Adaptive mode. In this mode, data is deduplicated while the backup process is going on. In Quantum's approach, the data stream coming into the DXi system during the backup process is written first to a disk buffer in small segments and then is deduplicated immediately. While the backup is still going on, unique blocks are being identified and added dynamically to the pool of unique blocks. When replication is enabled for a share or partition using adaptive mode, unique blocks are also being sent to the replication target at the same time. Buffering the data stream provides an important distinction from conventional in-line deduplication approaches. The buffer is designed to enable the system to adapt to input bursts, maximizing ingest speeds to minimize the backup window. Deduplication during backup does not require that any space be reserved to hold unprocessed backup sets.

The Adaptive mode is best suited to the widest array of situations—when minimal use of disk space is the most important consideration, when backup jobs come in bursts, and whenever replication needs to occur as quickly as possible. Applications like e-mail, backup of virtual servers, and edge-to-core replication of backed up data from remote offices to centralized sites are all natural applications. Performance can be very high—The DXi7500 can operate at up to 1.8 TB/hour (500 MB/sec) in this mode, validated by ESG Lab later in this report.

Deduplication as a post process: Quantum calls this policy Deferred mode. Deduplication is turned off during all or part of the backup ingest period and large data sets are written first in native mode to disk. Deduplication is then carried out as a post process operation. Because deduplication is suspended during initial write operations, all the processing power of the system is directed toward raw ingest, providing higher performance and shorter backup windows. The Quantum approach to post processing provides a level of flexibility that many post processing systems do not offer—the ability to schedule the operation. Quantum’s approach allows users to decide when to start the process, letting them defer deduplication until a backup job is finished or choosing to start deduplication at a specific time even if the backup is not completed.

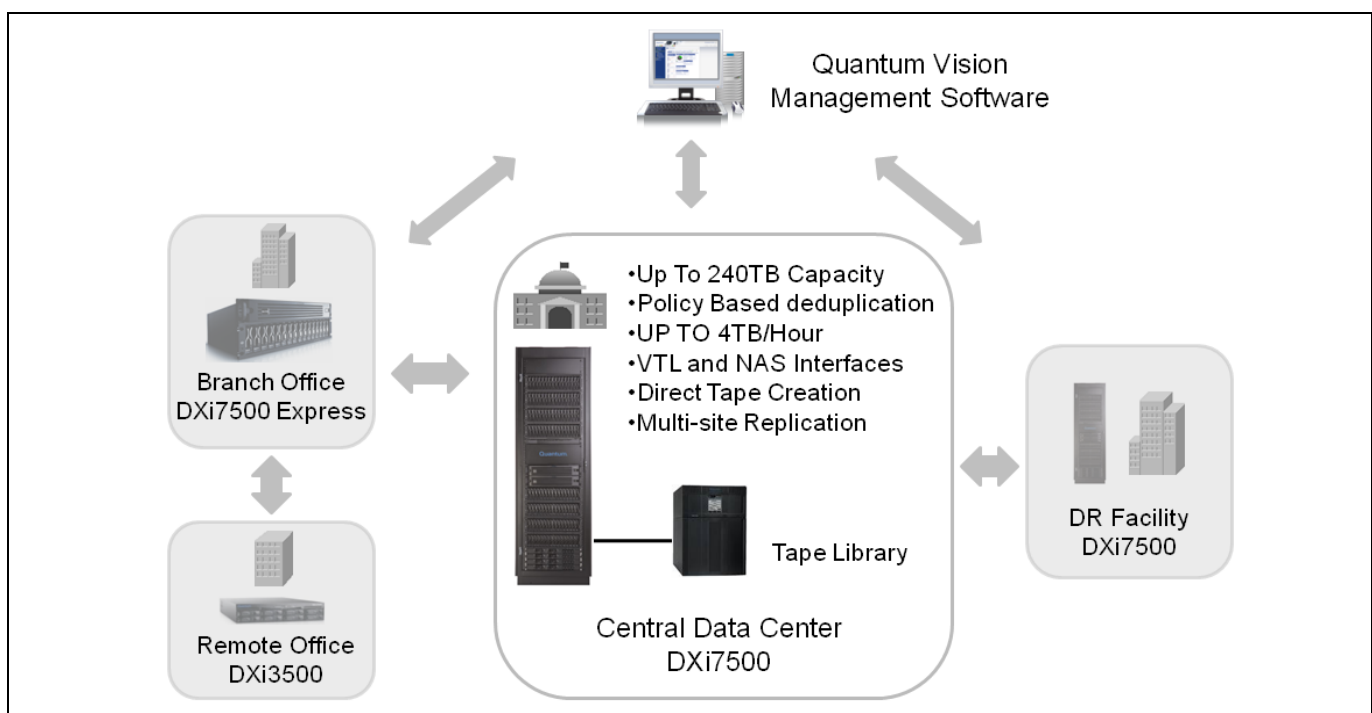
The deferred, or post-processing mode, is designed for situations that need the shortest backup window for deduplicated data and are able to reserve a disk landing area. Typical applications include online transaction processing (OLTP) systems, active databases, and other business critical applications that must have server operations back to full operating power as rapidly as possible.

Native mode operation: Deduplication is disabled. Although not strictly a deduplication mode, Quantum’s DXi7500 allows the use of some or all of the system in a conventional VTL or NAS mode. Data is stored either in native form or using standard compression algorithms without applying deduplication.

Native mode operations make the most sense for data that either does not benefit from deduplication or that will not be retained. Pre-compressed image files or encrypted data sets fit into the first category—pre-processing data normally makes it difficult to find repeated patterns and severely reduces deduplication effectiveness. Database recovery log backups fit into the second—these files are important to protect, but once the database rolls to the next point in time, they are no longer needed.

The DXi Series products all use the same deduplication technology and can be linked together using replication to help users create an enterprise wide edge-to-core backup architecture as depicted in Figure 3. Like other Quantum disk and tape automation systems, the DXi7500 is supported by Quantum Vision, a consolidated administration tool for managing multiple Quantum products, monitoring remote replication, and providing trend analysis reports for all monitored systems and processes. In addition, the DXi7500 supports the automatic creation of physical tapes through an integrated process that bypasses the application server.

FIGURE 3. QUANTUM DXI EDGE-TO-CORE DATA PROTECTION



This report explores the features and capabilities of the DXi7500 disk-based backup and replication system. ESG Lab tested the DXi7500's policy-based data deduplication technology, integrated tape creation capability, WAN replication, and user interface. Quantum Vision, the management system provided for disk, tape, and replication operations, was also examined.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of the DXi7500 disk backup and replication solution at a Quantum facility in Englewood, Colorado. DXi Series products integrate with existing backup software and policies as either a Virtual Tape Library or as a NAS disk device. When operating in virtual tape library mode, a DXi disk backup system emulates one or more tape libraries with both drives and cartridges that backup software utilizes exactly as it would physical libraries. When operating as a NAS device, the DXi systems present themselves as a CIFS share or NFS mount point. In this Lab Validation, ESG focused on the DXi7500's VTL capabilities.

Policy-based deduplication operation

To evaluate policy-based deduplication, ESG Lab backed up the same data set to a DXi7500 in each of the different modes to examine the relative differences in performance and operation using the three policies. Due to limitations in the quantity and power of the media servers available for test, these tests were not designed to test the absolute performance limits of the DXi7500.

FIGURE 4. THE ESG LAB TEST BED

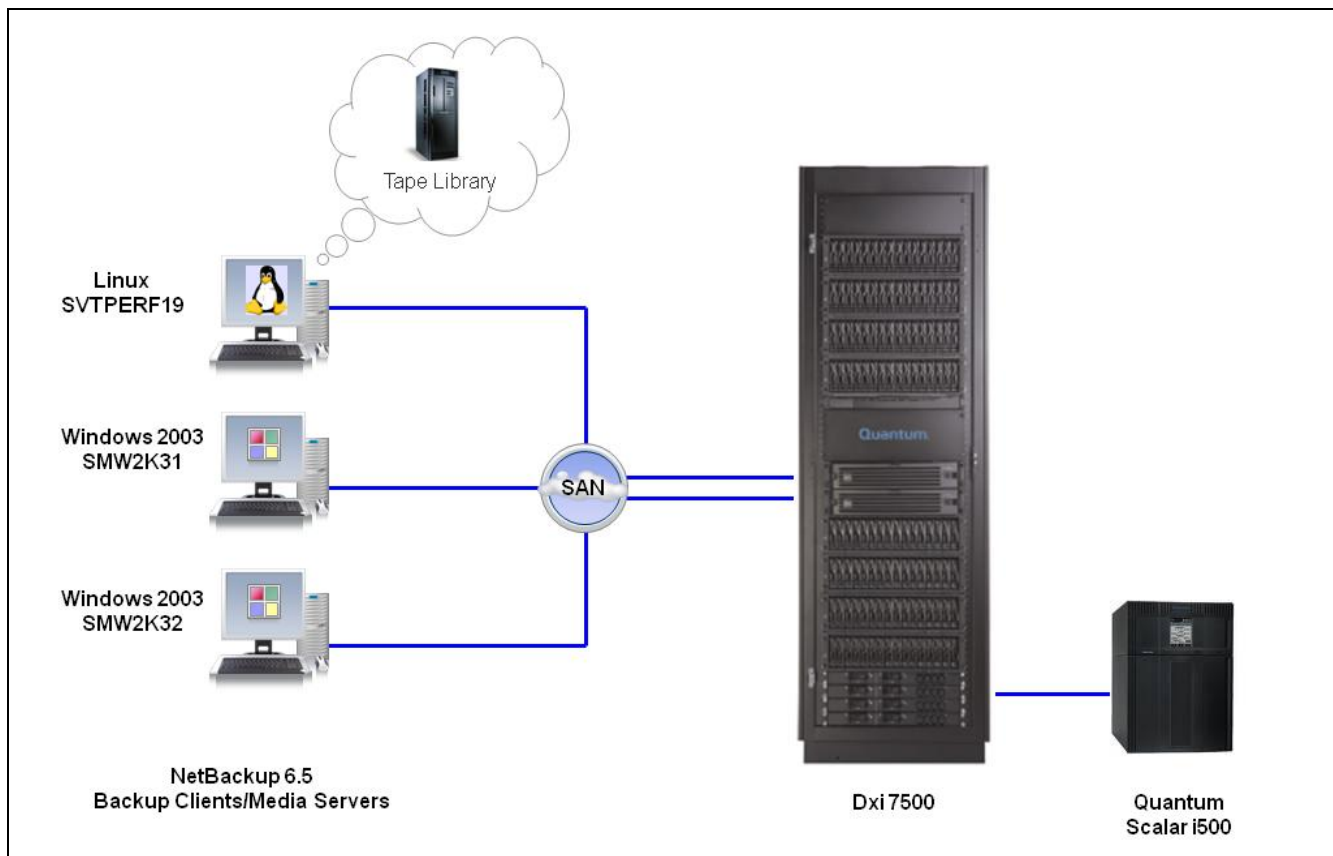


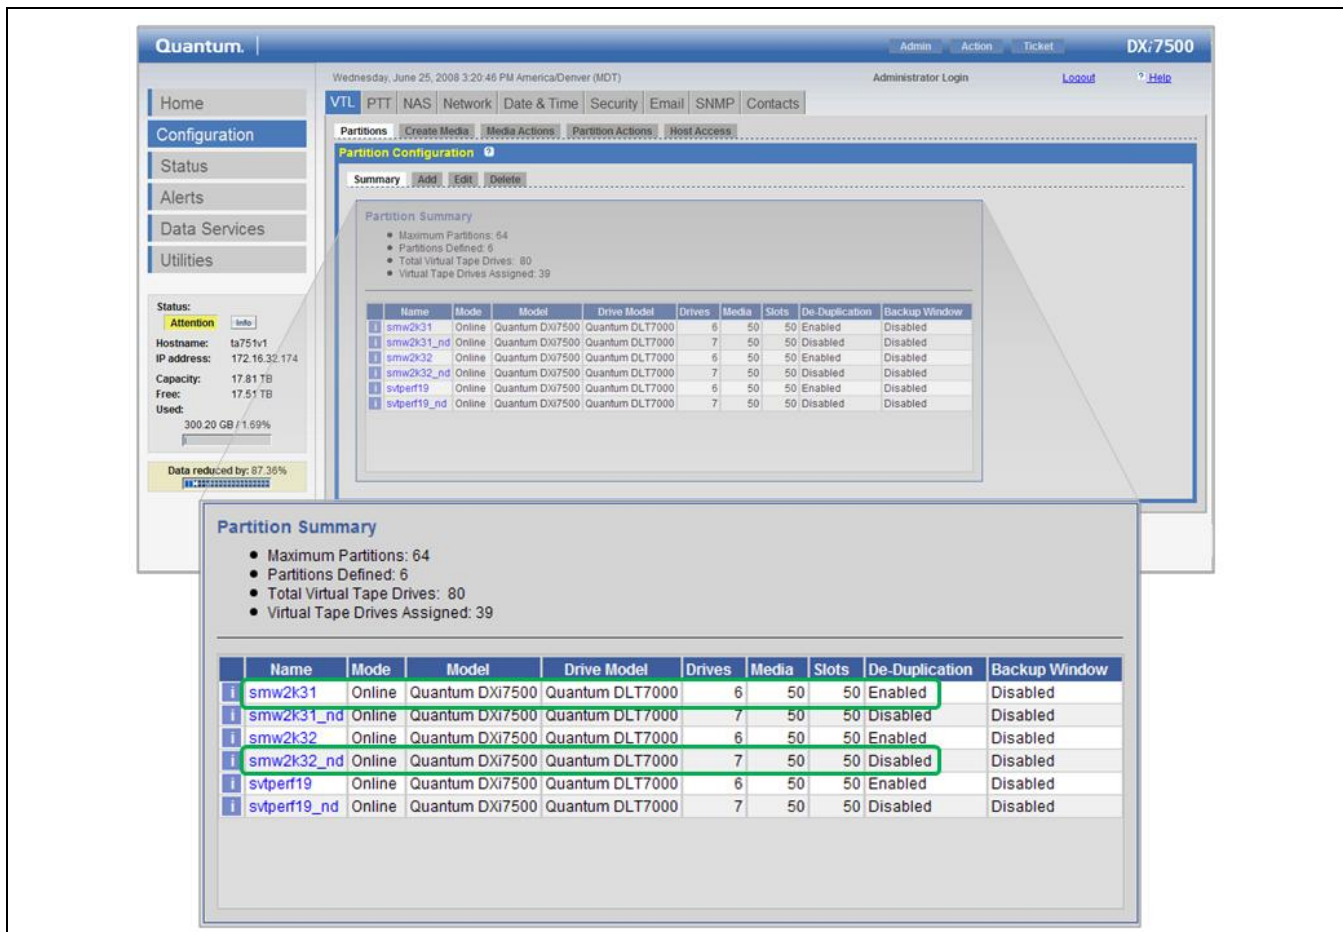
Figure 4 shows the test bed used in the hands-on testing described throughout this report. Testing began with a

pre-configured DXi7500—VTLs were pre-defined and virtual tapes had been created. ESG Lab used three servers: two running Microsoft Windows 2003 and one running RedHat Enterprise Linux configured as NetBackup media servers and backup clients. A Quantum Scalar i500 with two LTO3 drives was directly attached to the DXi7500 via Fibre Channel.¹

ESG Lab Testing

ESG Lab logged into the DXi7500 console and verified that VTLs were pre-configured in the system; half with deduplication enabled and the other half disabled. Each backup target, whether VTL or NAS, is called a partition in DXi terminology and deduplication is enabled on a partition by partition basis. Figure 5 shows the DXi7500 web console with the Partition Configuration summary highlighted. Status and configuration detail for each partition is displayed on this screen, including the number of virtual drives, slots, and tapes as well as whether deduplication is enabled for each partition.

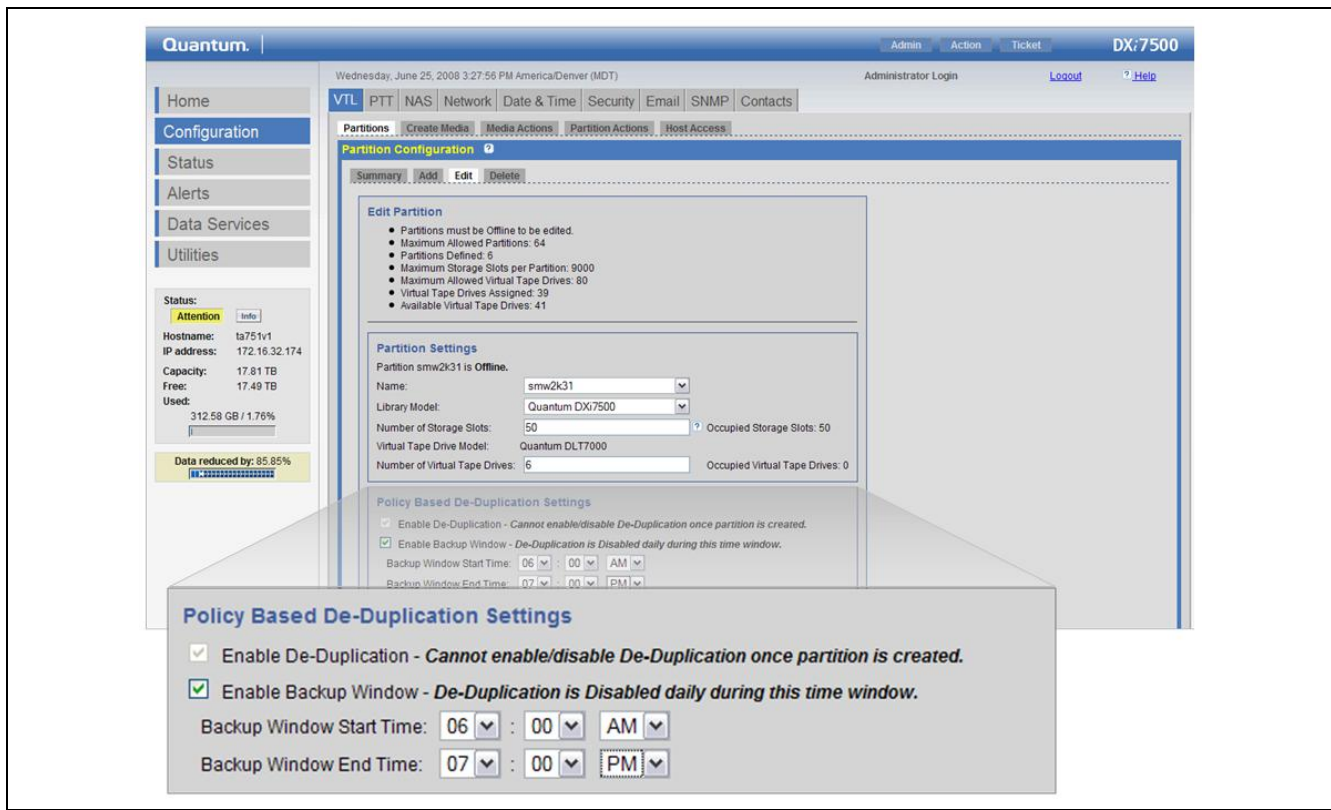
FIGURE 5. POLICY-BASED DEDUPLICATION



When creating a VTL for deduplication, the system defaults to Adaptive mode. To set a partition to Deferred mode, the administrator needs to set a 'backup window,' which is the times between which deduplication will not occur. Figure 6 shows the setting for one partition under test. A customer would typically set this window to match the backup window for the systems being backed up to this partition, allowing backups to occur without any impact from deduplication. As seen in Figure 6, ESG Lab set the window for 6AM to 7PM to ensure that deduplication for this partition would not occur during testing, but overnight, while the system was idle.

¹ Configuration details can be found in the Appendix.

FIGURE 6. CONFIGURING DEFERRED DEDUPLICATION



ESG Lab performed backups of twelve consecutive days worth of real world MS Exchange data provided by Quantum IT to test the ability of policy-based deduplication to back up multiple data sets using adaptive, deferred, and native policies. Backups were configured and run from within NetBackup using exactly the same methods as with physical tape libraries.

Why This Matters

ESG research² indicates that cost and impact to backup and restore performance are two of the leading obstacles to disk-based data deduplicated backup deployments. Policy-based data deduplication addresses both of these issues by reducing the amount of data retained on disk while providing optimal performance for different data types as well as backup and retention policies.

ESG Lab has validated that Quantum policy-based data deduplication can be used to reduce disk capacity by up to 98% depending on the type of data, the change rate between backups, the backup policies in place, and the retention period applied. ESG Lab tested all three of the DXi's policy-based deduplication modes and found that by applying policies appropriate to data type and corporate policy, customers can effectively provide high performance backup services, plus fast and reliable restores using greatly reduced disk capacity. This lowers the cost per GB for backup data and enables companies to retain data exponentially longer for recovery purposes while minimizing the impact of deduplication on backup windows and recovery SLAs.

² Source: ESG Research Report, *Data Protection Trends*, 2008

Policy-based deduplication performance

A tape drive can only perform one backup at a time. To get more than one backup job running at the same time, more tape drives need to be added and run in parallel. A disk-based backup and recovery solution with many random access disk drives emulating many virtual tape drives (up to 64 for the DXi Series) can run many backup jobs simultaneously. Much like the difference between a DVD and a VHS tape, the random access of disk also provides improved performance when locating files to be restored.

ESG Lab Testing

ESG Lab performed a series of backups using both adaptive and deferred deduplication, as well as native VTL mode (no deduplication), to examine the relative performance of a DXi7500 in each mode. Nightly full backups were simulated with data sets collected from a production Microsoft Exchange environment. Veritas NetBackup was used to perform backups. ESG Lab began with one server backing up one stream and gradually ramped up the number of concurrent servers and streams to three servers and eighteen streams.

FIGURE 7. DXi7500 PERFORMANCE COMPARISON ACROSS DEDUPLICATION MODES




Figure 7 shows the first set of backups performed, starting with one server running one stream and scaling up to three servers running three streams each. As expected, Native VTL mode, where no deduplication is taking place, shows the highest relative performance, while Adaptive deduplication shows the greatest impact to backup performance.

ESG Lab also audited Quantum's in-house testing to validate the maximum performance claimed by Quantum for the DXi7500. Quantum used a 72 TB (usable capacity) system, which is less than half of the maximum capacity of a DXi7500i, to conduct these tests. The testing environment and data sets were similar to those used by ESG Lab. The DX7500i was configured as a Fibre Channel VTL in each of the three policy-based deduplication modes (Adaptive de-dupe, Deferred de-dupe and Native mode with compression). A total of four powerful host servers running 16 streams were used for testing (four streams per host) and each data set was used as a stream to a virtual device on the DXi7500.³

Microsoft Exchange data was used as the source for backups and Microsoft Exchange Server Load Simulator (LoadSim) was employed as a benchmarking and deployment verification tool to simulate the delivery of multiple messaging requests from MAPI clients to a server running Exchange. The backup software and performance results were reported through Symantec NetBackup 6.5 and internal logging tools. Test results obtained by ESG Lab and Quantum are shown in Table 1.

TABLE 1. DXI 7500 PERFORMANCE RESULTS – ESG LAB TESTED AND AUDITED


ESG lab Tested - 18 TB Usable		
Deduplication Policy	Multi-stream Ingest Performance	
	MB/sec	TB/Hour
Adaptive deduplication mode	263	.9
Adaptive deduplication mode – 2 nd full	297	1.1
Deferred deduplication mode	467	1.7
Native VTL Mode with compression	555	1.9
ESG lab Audited - 72TB Usable		
Deduplication Policy	Multi-stream Ingest Performance	
	MB/sec	TB/Hour
Adaptive deduplication mode	500	1.8
Deferred deduplication mode	889	3.2
Native VTL Mode with compression	1111	4.0

What the Numbers Mean

- Performance scaled smoothly in all modes as the number of servers and streams were increased.
- Native VTL mode offered the highest performance, followed by Deferred, then Adaptive mode.
- The difference in performance between each of the methods was negligible at low traffic levels. ESG Lab believes that this indicates that most organizations can take advantage of the capacity savings of Quantum DXi7500 deduplication without a performance penalty.
- The DXi7500 configuration tested by ESG Lab was able to back up between one and two terabytes of data per hour depending on the deduplication policy.
- ESG Lab noted an 11% increase in performance between the first and second Adaptive mode deduplicated full backups. Subsequent deduplications of full backups have less new data to process, but still must compare the full backup stream to previously deduplicated data.
- The DXi7500 was able to demonstrate enterprise-class backup performance in an active Microsoft Exchange environment typical of many large enterprises.
- The larger number of more powerful servers used in Quantum's testing were able to drive more than twice the IO from the DXi7500 than ESG's hands-on tests.

³ Configuration details can be found in the Appendix.

FIGURE 8. ESG LAB TESTED AND AUDITED PERFORMANCE RESULTS COMPARED



As can be seen in Figure 8, the DXi 7500 tested by Quantum under extreme load showed similar performance differences between deduplication modes as the DXi7500 tested by ESG Lab.

Why This Matters

ESG research⁴ has found that the number one challenge enterprises report with their data protection processes and technologies is the need to reduce backup times. Backup administrators have been struggling for years to get nightly backups completed before business resumes in the morning. Quicker recoveries are also cited as needed to increase user productivity and meet service level agreements.

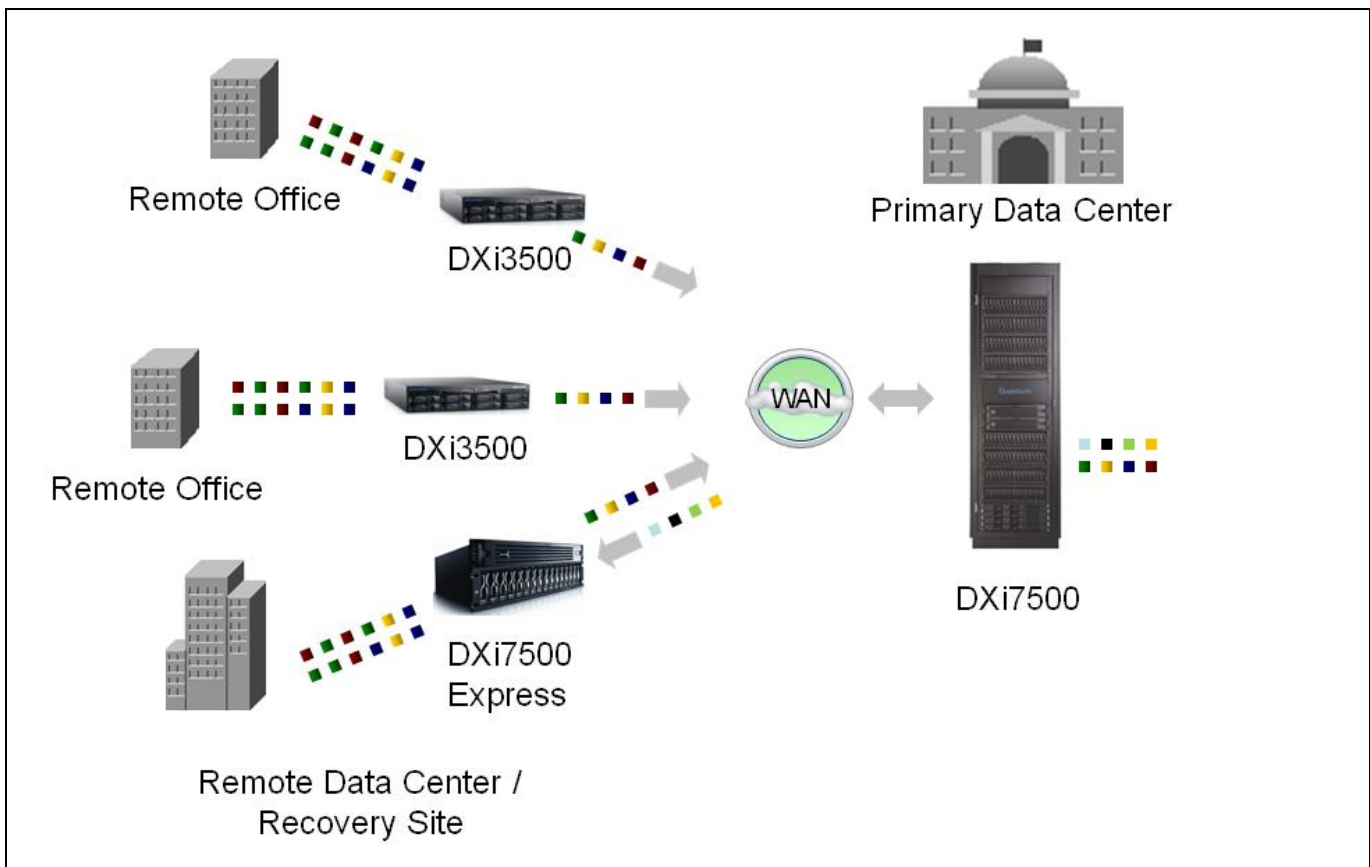
ESG Lab validated through direct test and audit that Quantum's DXi7500 can easily sustain its specified aggregate backup throughput ratings as advertised. In other words, a single Quantum DXi7500 disk backup system can be used to protect up to 32 TB of data in an eight hour shift (depending on deduplication method) and restore individual files in a matter of seconds. The ability to choose between multiple deduplication modes in the DXi Series means that users can meet the protection requirements of a large number of different data types with one system, enabling optimal performance for each data set in an enterprise while lowering acquisition costs and operational complexity.

⁴ Source: ESG Research Report, *Data Protection Trends*, 2008

Remote Replication

Backup data residing on a DXi Series disk backup solution can be replicated over a wide area network for secure, network-based disaster recovery protection. Data deduplication reduces the amount of data that needs to travel over the WAN, enabling a global backup replication strategy and reducing dependency on physical tape transport for disaster recovery. With the DXi Series, replication occurs automatically without waiting for the backup jobs to complete. Data traveling across the WAN is encrypted via 256 bit AES. Many to one and bi-directional replication configurations are supported as shown in Figure 9. In this example, DXi Series appliances at multiple remote offices are configured to replicate automatically to a DXi7500 at a corporate data center. The DXi7500 at the primary data center is replicating critical application data to one of the remote sites for disaster recovery protection.

FIGURE 9. REMOTE REPLICATION



ESG Lab Testing

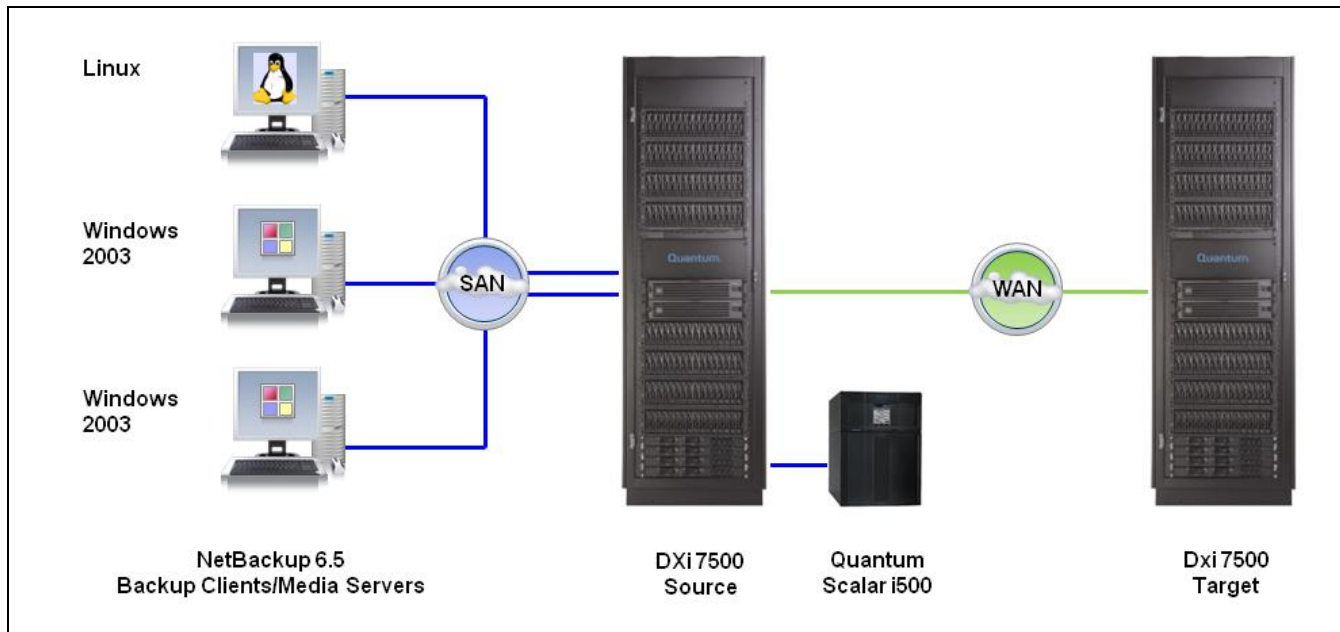
ESG Lab configured remote replication between two DXi7500s located in Englewood, Colorado as shown in Figure 10. The DXi7500 web console was used to configure and monitor replication over LAN connection.

ESG Lab first enabled replication on a VTL partition in the source DXi7500, TA751, and configured the target (TA752) to accept replication from the source. Next, a full backup of a directory was initiated by ESG Lab. Replication of unique blocks started automatically, as soon as data began to be ingested, and it continued during the backup as a background process.

In addition to copying the actual unique blocks, the replication process also sends namespace data (metadata about the virtual cartridges or the NAS backup files) to synchronize the data on the target with the data on the source. Namespace data can be replicated daily based on a schedule or on-demand. In this test, ESG Lab

replicated the namespace on-demand to ensure that the namespace data was replicated before failover and failback. Clients would usually schedule namespace replication to occur after all nightly backup jobs have completed.

FIGURE 10. TESTING DXi7500 REMOTE REPLICATION



To simulate an outage at the primary data center, ESG Lab deleted the source VTL partition from the source DXi7500 and failed over to the DXi7500 target. ESG Lab then recovered the partition on the target system and restored the dataset to an alternate location. Validity of the restored data was verified by physical inspection of files in the new destination and running compares against the original source files. File name, size, and content matched the source files exactly.

Finally ESG Lab reversed replication, failed back to the source-side DXi7500 and confirmed that the VTL was completely restored. A final verification of the contents and integrity of the replicated backup data was performed using Veritas NetBackup. To provide an additional level of protection and security, the target DXi7500 holds multiple snapshot-style copies of previous states, allowing users to preserve multiple known good images of the backup replica.

Why This Matters


Because of the large capacity of backup archives and the high cost of WAN bandwidth, traditional backup to disk solutions prevent customers from replicating data to a remote site, forcing them to rely on tape-based data transport. Assuming the media is good, finding and restoring applications from tape after a disaster is difficult at best.

ESG Lab has verified that Quantum DXi Series data deduplication brings the same level of compelling economic value to WAN bandwidth optimization as it does to disk capacity. Customers can essentially reduce their bandwidth requirements up to 50 times, enabling a global data protection strategy for data centers and remote offices. Reducing bandwidth not only reduces cost but also can be used to protect enormous amounts of data over long distances, which otherwise may have been impossible.

Path-to-tape

In most backup environments, physical tape is still used for long-term and offsite archiving. Organizations that have completely embraced backup-to-disk and VTL technology still frequently need to export virtual tapes to physical tapes for long term retention, archiving, duplication or cloning as well as to use physical tapes to restore archived data. Traditionally, the backup application would need to use the Media Server to copy tapes, pulling all the tape data from the VTL and writing it out to a physical tape library, as seen in Figure 11.

FIGURE 11. EXPORTING TAPES TRADITIONALLY USING THE MEDIA SERVER



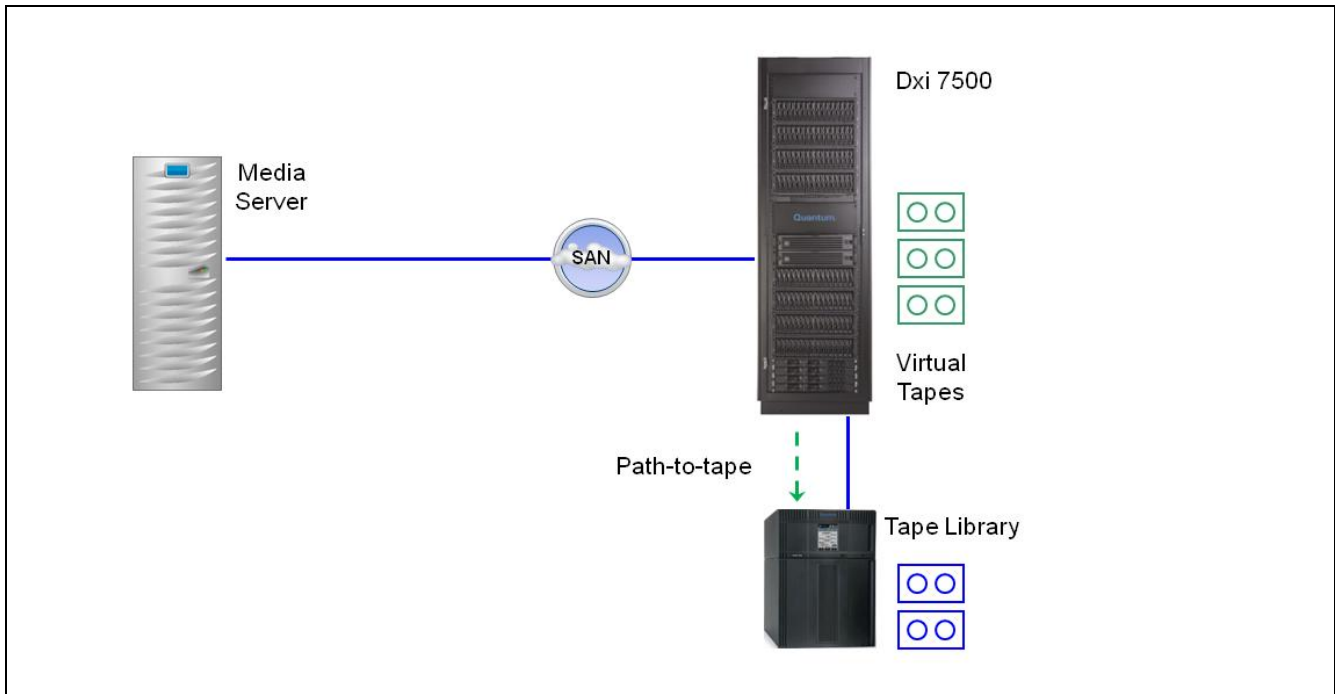
In this scenario, all backup data flows through the SAN from the VTL to the media server and then back across the SAN to the tape library, consuming significant resources on the media server and I/O on the SAN that could otherwise be used to perform backups and restores.

The DXi7500's integrated path-to-tape allows it to directly copy virtual tapes to a physical tape library without involving the media server in data movement. The DXi7500 offers IT the capability of integrating physical tape creation without using the backup application's media servers to move data while coordinating that movement with supported backup software. This allows the backup software to initiate the tape creation process, to track the different instances of the backup sets, and to manage them with different retention and expiration policies while moving the actual transfer process off the media server. Integrating tape creation in this manner gives users a multi-tier backup system that is fully integrated with the backup application's media management.

Since the backup application has full knowledge of the process, users can change tape technology (send data from a virtual SDLT cartridge to a physical LTO cartridge), and either span media (send one virtual tape to multiple physical tapes) or concatenate (include several virtual tapes on one piece of physical media). The application automatically tracks the location of the data across the two sets of media.

To make this work, the DXi uses the Network Data Management Protocol (NDMP) to manage the data transfer. NDMP is a protocol originally developed to transport data between NAS devices and backup devices, but it has subsequently been expanded to include different functions, including tape creation. In both cases—NAS transfer and tape creation—the protocol removes the need for transporting data through the backup server itself, enhancing performance and reducing the operational load on the media server. Figure 12 shows one possible configuration for path-to-tape, where the physical tape library is directly attached to the DXi7500. The tape library can also be SAN attached.

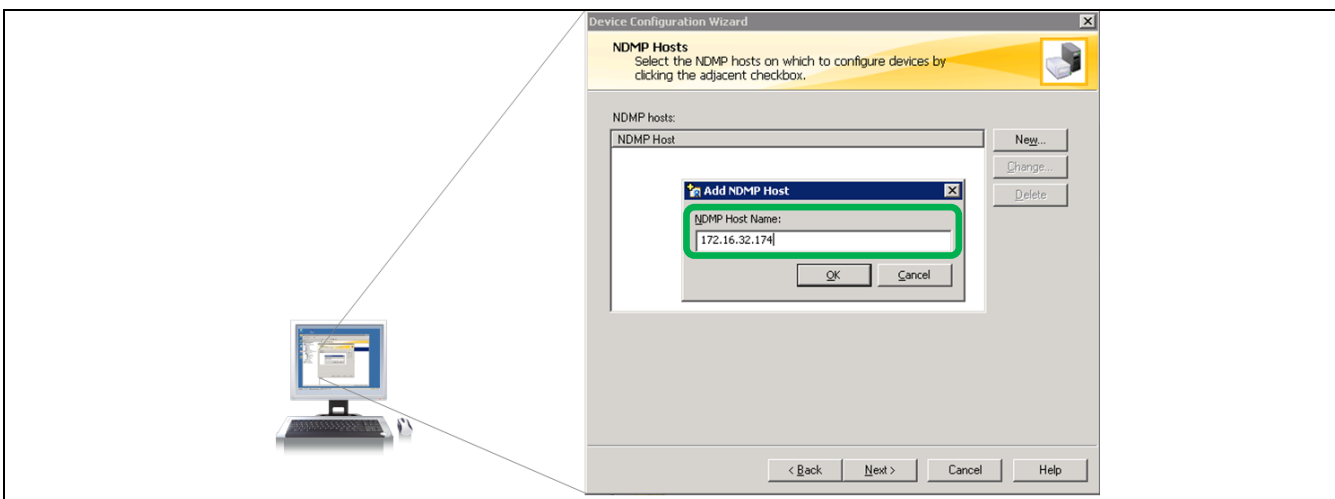
FIGURE 12. DIRECT TAPE CREATION WITH PATH-TO-TAPE



ESG Lab Testing

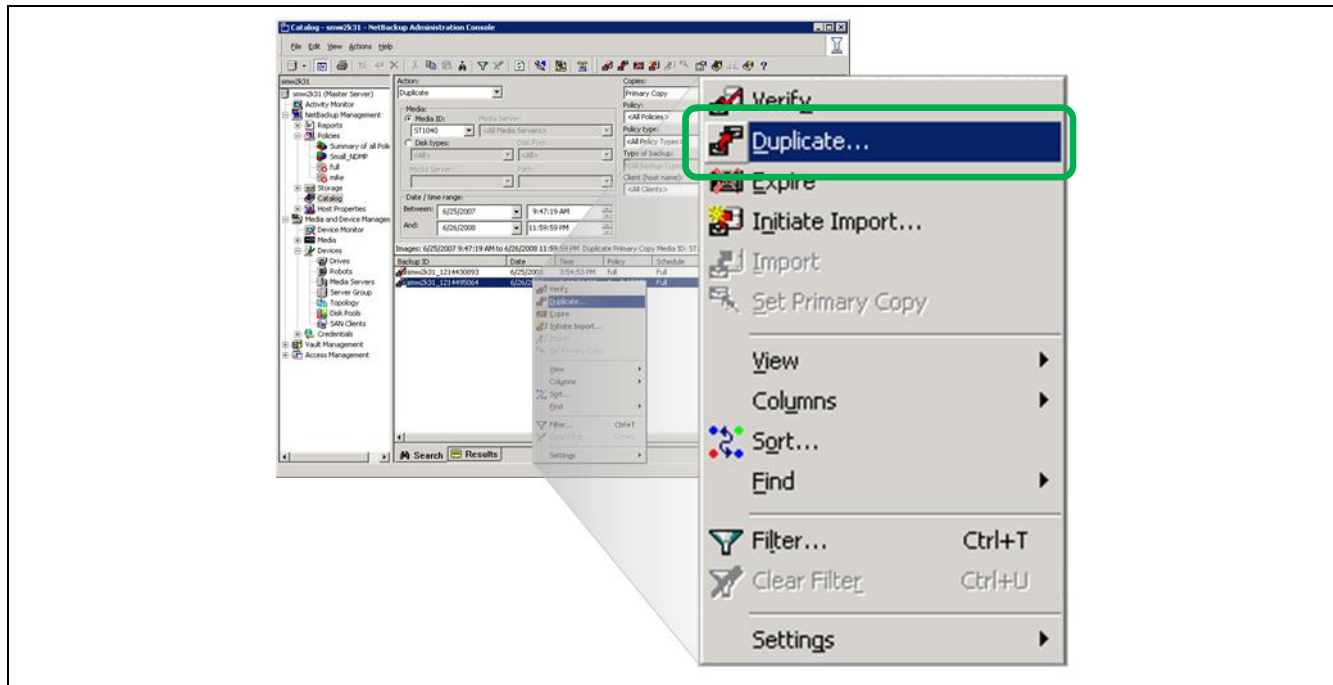
To test path-to-tape functionality, ESG Lab configured the DXi7500 as an NDMP host in NetBackup, as seen in Figure 13, by entering the IP address of the DXi7500 in the NetBackup device configuration Wizard.

FIGURE 13. CONFIGURING NETBACKUP FOR DIRECT TAPE CREATION



This allows the DXi7500 VTL to control the tape library directly when commands are issued from the NetBackup console. Next, a backup was run to the VTL, which created a virtual tape cartridge (VTC). ESG Lab then issued the Duplicate command against the VTC from within NetBackup as seen in Figure 14.

FIGURE 14. CONFIGURING NETBACKUP FOR DIRECT TAPE CREATION



The Duplicate command was passed to the DXi7500, which copied the backup data directly to physical tapes in the attached tape library. An important note: the physical tapes are written in native format—that is, they look like any other tape written by the backup application and they can be read directly on any compatible physical tape drive. This allows the files to be read directly back to the server during restores—there is no requirement to use the disk-based system during restores from physical tape.

Once the copy was complete, ESG Lab marked the physical tape as the primary copy in the NetBackup catalog. ESG Lab next performed a restore of the backup set from the physical tape to an alternate location. Finally, ESG Lab verified the validity of the restored files by physical inspection and by comparing the restored files against the original source files. The file names, sizes, and content matched the original source files exactly.

Why This Matters

ESG research⁵ shows that 87% of organizations integrating disk-based VTL into their backup strategy still incorporate physical tape in their backup policies. Tape copies are made for offsite DR, archival, and compliance requirements. A disk-based backup solution's integration with physical tape enables IT managers to meet these requirements with existing legacy tape systems and to significantly downsize their tape infrastructure as they move an increasing amount of operational backup and restore operations onto disk-based systems.

ESG Lab validated that Quantum DXi path-to-tape functionality is fully compatible and integrated with backup applications' built-in media management functions. Path-to-tape moved data directly from the DXi7500 to tape without using the backup application's server to move data. This reduced media server loads and SAN traffic, shortening the time required for media creation and allowing media servers to perform other tasks more effectively. ESG lab controlled the process using the backup application's media management system which provided fully independent control of disk and tape tiers. This allowed different expiration policies for the same data in different locations, all under the same single point of management used today for expiration policies. These capabilities provide effective control, lowered costs, and reduced management overhead.

⁵ Source: ESG Research Report, *Data Protection Trends*, 2008

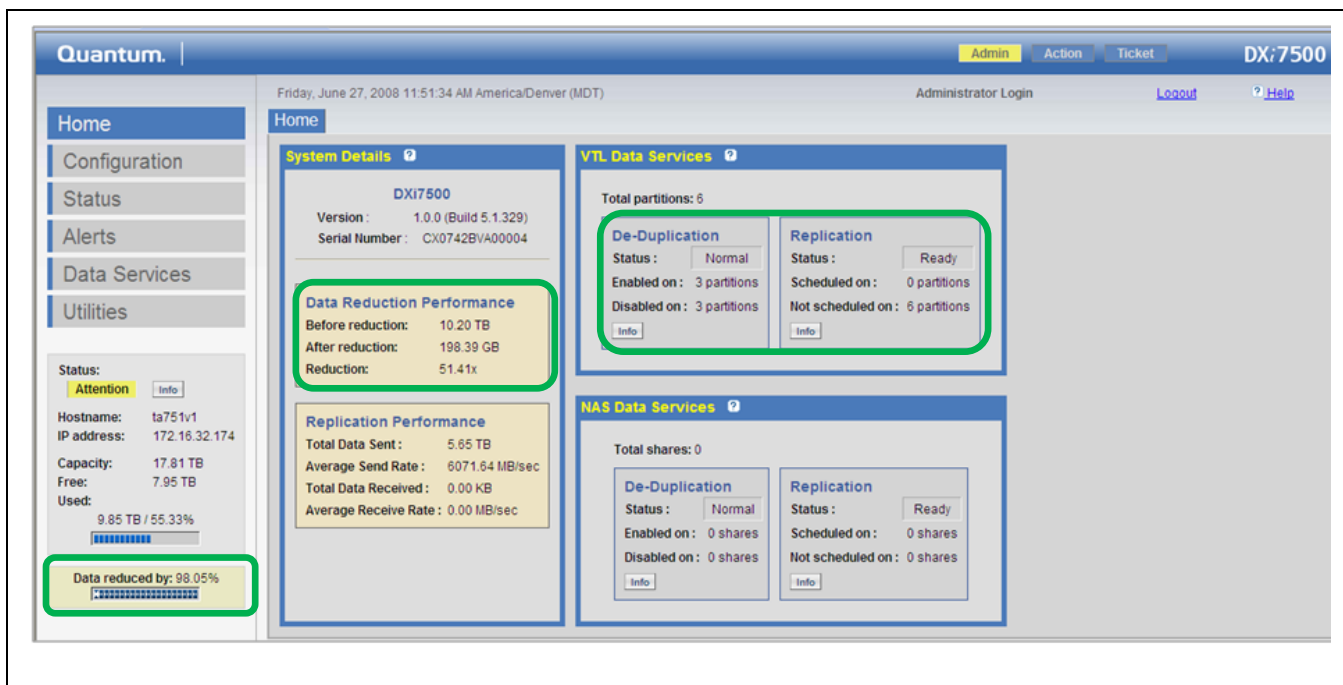
Manageability

One of the major advances that deduplication technology has enabled is allowing users to link multiple sites together in an edge to core strategy using remote replication and integrating disk and tape tiers into a single, consolidated backup architecture. As backup systems become more interconnected and integrated, managing the system becomes a critical element of any solution set. This section will examine the management features of the DXi7500 and the capabilities of Quantum Vision, a management software tool that allows users to manage disk, tape, and replication in multiple sites from a single console. The report will also discuss Quantum's service organization, consultative services, and the company's remote monitoring and diagnostic tools that support them.

ESG Lab Testing

ESG Lab found the DXi7500 GUI to be easy to navigate, with plenty of useful information right on the home screen. Administrators can examine data deduplication performance for the entire system at a glance, expressed as both a ratio and as a percentage as displayed in Figure 15.

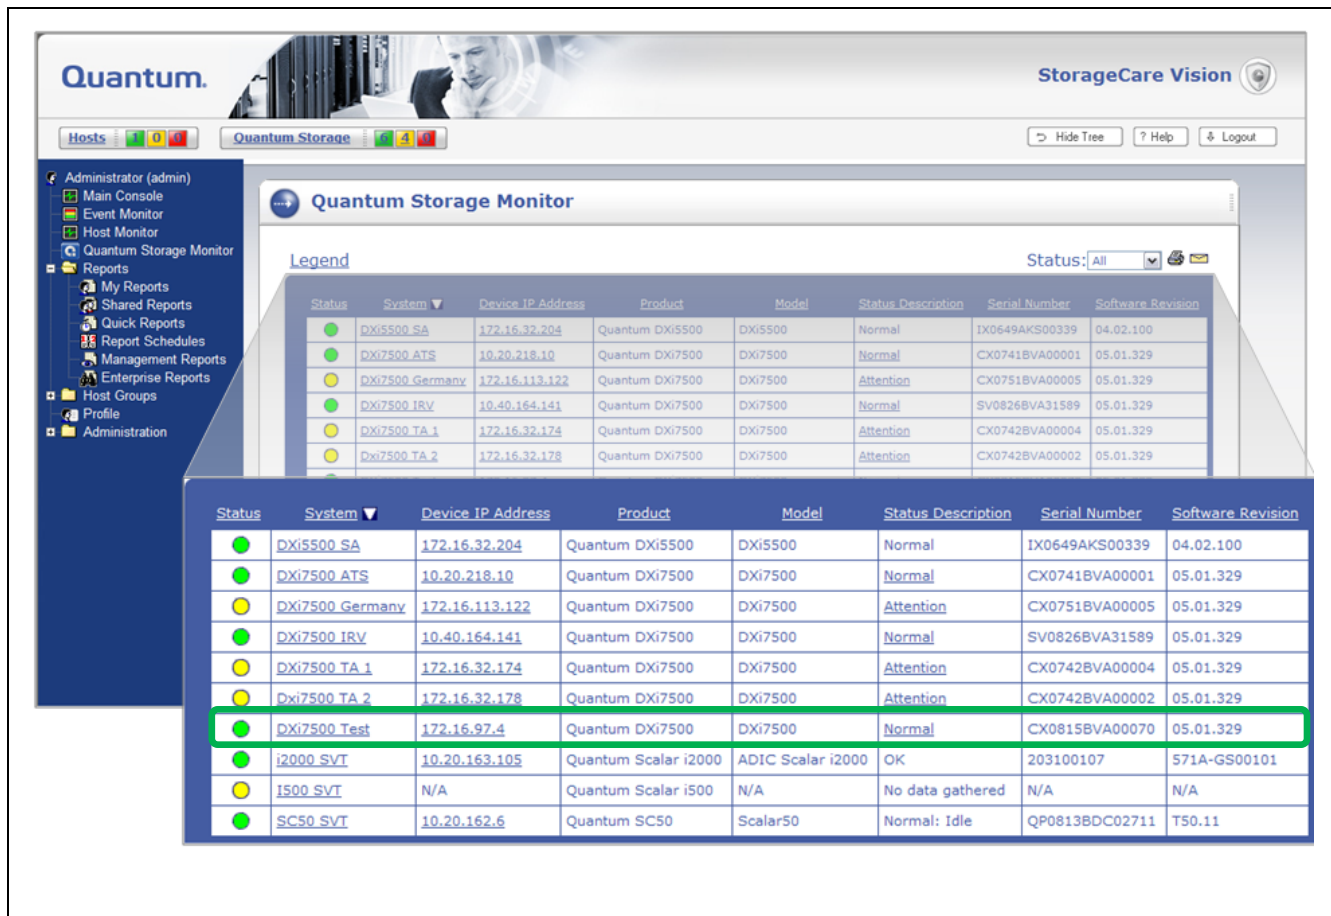
FIGURE 15. THE QUANTUM DXI7500 GUI



Also displayed on the home screen, the status of data deduplication and replication give administrators a quick status check of the overall health of the system.

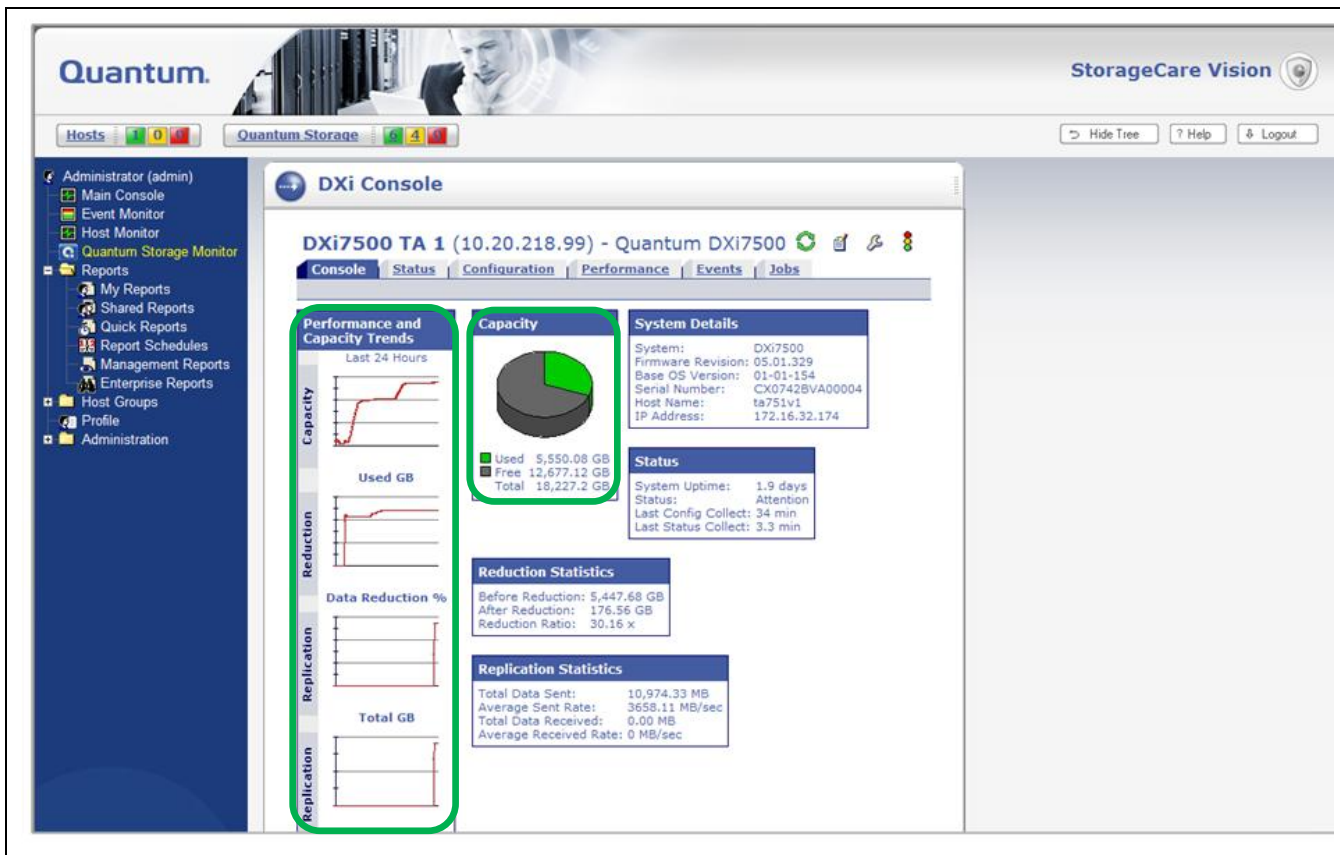
ESG Lab next examined Quantum Vision, designed to provide a complete view of the backup environment, including tape and disk storage assets and the process of replicating data between disk systems. When ESG Lab logged into the Quantum Vision web console, the Quantum Storage Monitor presented a list of all systems under management. This global view provided a snapshot of the status and health of every backup device in this petabyte-plus environment with the ability to filter by status and quickly drill down into any system that required attention as shown in Figure 16.

FIGURE 16. QUANTUM VISION



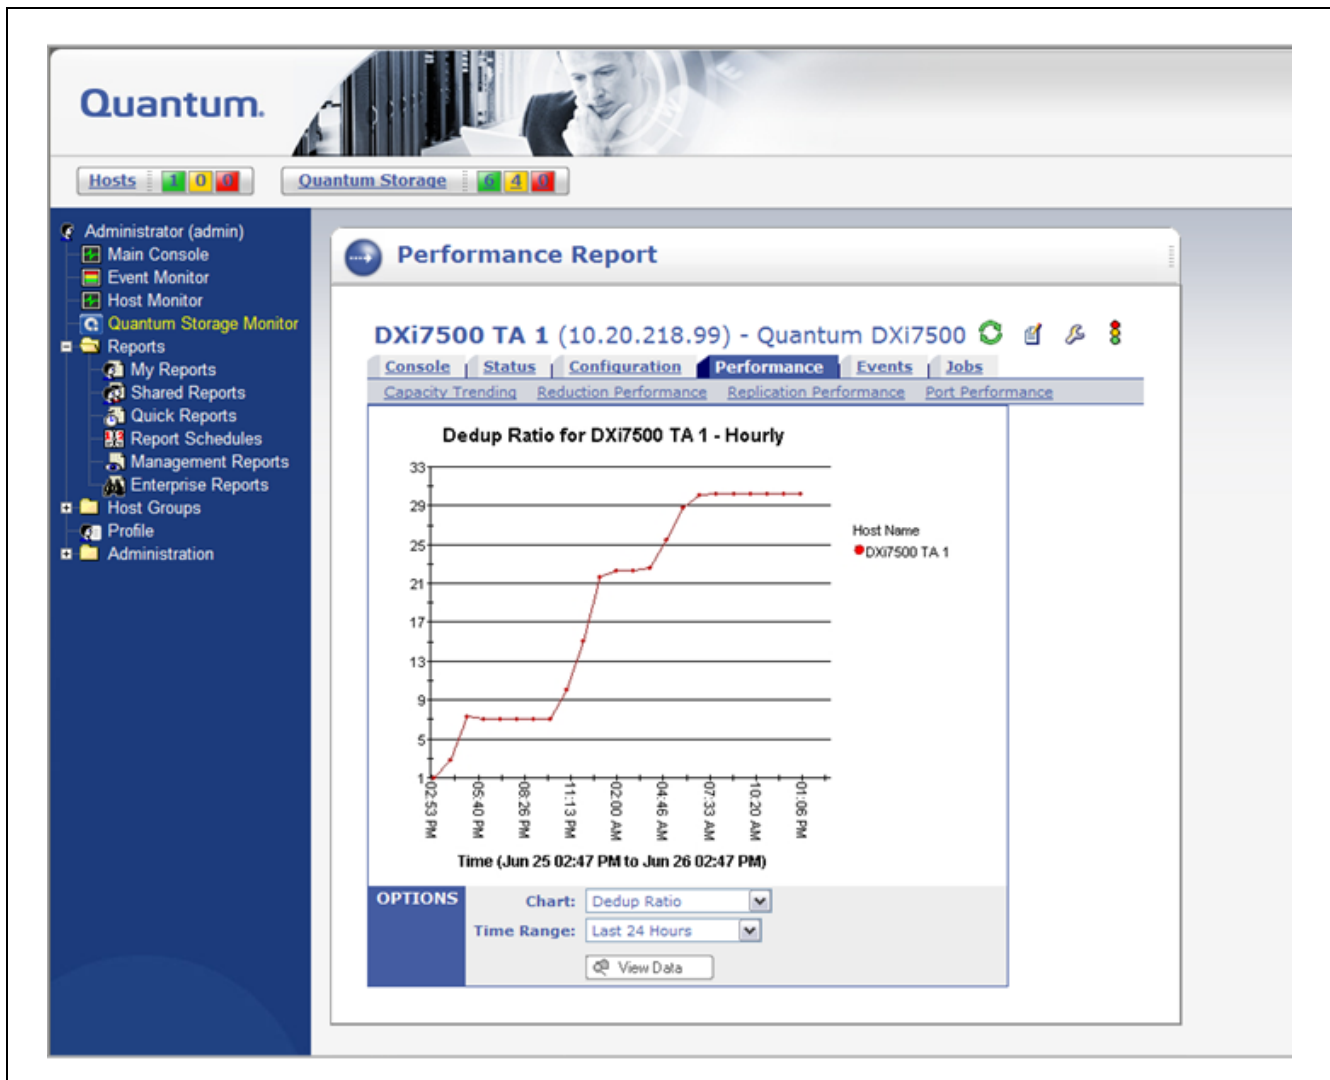
Next, ESG Lab examined the DXi7500 Console within Quantum Vision. Several graphic representations of performance, capacity, and data reduction are offered to administrators, as seen in Figure 17. If replication is enabled, users can also look at status of replication jobs and monitor replication performance trends.

FIGURE 17. QUANTUM VISION DXI CONSOLE



In addition to the expected configuration and management functionality, ESG Lab found Quantum Vision's reporting and alerting functionality to be extensive and robust. Administrators can custom-design reports showing history and trending for capacity, performance, and deduplication and have those reports e-mailed out to multiple addresses. The reports, like the deduplication performance report seen in Figure 18, are clear, concise, and easy to read. They provide both point in time views and trend analysis reporting for any combination of units anywhere in the world. This report shows the overall deduplication ratio climbing as each successive full backup was completed.

FIGURE 18. QUANTUM VISION DEDUPLICATION PERFORMANCE REPORT



Why This Matters

Ease of implementation and management was reported to ESG as one of the top concerns when choosing a data deduplication backup solution. This will become more important as users leverage the technology to create integrated enterprise-wide backup environments that link systems via replication and integrated tape creation. A data deduplication disk backup solution that is easy to manage and integrate into existing environments saves time, bandwidth, and money. Customers choose a solution with the expectation of robust management for the data center as well as remote locations where IT resources may be limited or non-existent. In the case of Quantum's DXi Series, ESG Lab was not disappointed. Configuring the DXi7500 is straightforward and intuitive. Ten minutes after beginning the configuration, Veritas NetBackup jobs were storing backup data on a Quantum DXi7500. ESG Lab also liked the look and feel of the Quantum Vision application, which provided excellent management, monitoring, and capacity and performance trending for replication operations as well as disk and tape devices.

ESG Lab Validation Highlights

- ☑ After the DXi7500 had been configured as a virtual tape library, Veritas NetBackup worked exactly as if it were a real tape library. Configuring deduplication policies and schedules on the DXi7500 was straightforward and intuitive. ESG Lab appreciated the ability to define three different types of policies to meet a wide variety of performance and capacity requirements: Adaptive, Deferred, and Native.
- ☑ Multiple single and multi-stream backups were performed using Veritas NetBackup along with an audit of scale-out tests performed by Quantum. ESG Lab verified Quantum's claims of backup performance up to 4 TB per hour for the DXi7500
- ☑ Path-to-tape integrated smoothly with NetBackup media management and offloaded resource intensive copy functions from the media server while using familiar NetBackup commands.
- ☑ Remote replication was configured to simulate a site failure and disaster recovery scenario. ESG Lab was able to fail over to a DR site and begin restoring files in less than a minute. Failback to the primary was also easy and just as fast.
- ☑ Quantum Vision was visually appealing and provided excellent management, monitoring, and capacity and performance trending.

Issues to Consider

- ☑ As with all VTL's today, when a cartridge is deleted or expired in a backup application, space on the DXi is not reclaimed until the cartridge is deleted, expired, or overwritten via the DXi management application and the Reclamation process is run. Tighter integration with Backup applications to automatically trigger a delete or expire in the DXi when a cartridge is expired in the backup application would be a useful enhancement.
- ☑ The DXi has clearly moved squarely into the enterprise with the capacities supported in the DXi7500, leveraging the StorNext file system for improved high availability. Using a multi-node, clustered design would enable even larger centrally managed pools of deduplicated capacity.
- ☑ Active Directory integration for Vision would simplify role-based administration and management for a global storage and backup environment leveraging an infrastructure already present in most enterprises.

ESG Lab's View

ESG Lab has been tracking the progress of Quantum's disk-based backup and recovery solutions for several years, conducting its first hands-on testing of the company's DX100 enterprise VTL system in April, 2005 and validating DXi Series appliances in 2007. ESG Lab testing and discussions with end-users have confirmed that Quantum's disk-based backup solutions fit seamlessly into existing backup environments while providing dramatic performance and capacity reduction benefits compared to legacy tape-based methods. The DXi7500 enables a comprehensive, enterprise wide edge-to-core data protection architecture that goes beyond disk-based backup.

During this third independent lab validation, ESG Lab confirmed Quantum's move into the enterprise space with the DXi7500. Quantum's single, scalable architecture across the DXi series provides flexible policy-based data deduplication technology delivering dramatic disk capacity savings while offering scalable, predictable performance that is optimized for both the type and value of data.

A modest configuration tested by ESG Lab was able to back up at a sustained 2 TB/hour while at the high end, Quantum was able to demonstrate nearly 4 TB/hour. Quantum's bandwidth-optimized remote replication was able to reduce network bandwidth significantly when replicating backup data over the WAN, enabling an edge-to-core data protection strategy covering remote and branch offices as well as multiple data centers. Path-to-tape integration with backup software enables enterprises to meet offsite and deep archive requirements using the familiar tools and techniques of their backup application without burdening media servers or consuming SAN bandwidth with the task of tape creation.

Quantum acquired ADIC, including the foundational variable length block-level data deduplication technology ADIC had previously acquired from Rocksoft in late August, 2006. ESG was pleasantly surprised to learn that this patented data deduplication technology had been built into the DXi Series platform when it was announced only a few months later in December, 2006. What was accomplished in a very short time is very impressive. Since then, Quantum has demonstrated it can leverage its significant customer-facing resources, global services and support, and huge tape library install base to expand its presence in disk and software into the enterprise. Going forward, ESG Lab believes that the DXi Series provides Quantum and its customers with an excellent foundation for the deployment of advanced data protection solutions.

Quantum's Vision software exemplifies the company's depth and breadth of end-to-end solutions for backup, recovery, and archive encompassing disk and tape. It should bring significant value to customers grappling with the challenges associated with effective management of their data protection resources. ESG Lab found tools such as the deduplication trend analysis to be an excellent example of the deep, yet intuitive value of Vision's rich set of capabilities.

ESG Lab believes that the combination of policy-based data deduplication, enterprise class performance, and scalability—along with comprehensive storage management software and services—provides a unique approach for optimizing data protection and recovery strategy in the enterprise. Quantum customers can now retain more data for fast and reliable restores and longer retention periods while minimizing impact on backups with policy-based deduplication. Combined with the enhanced capacity, performance, and flexibility of the DXi7500 platform, customers have a wide choice of configurations in a single solution, which can be used to dramatically increase the role of disk in the protection of critical data.

Appendix

TABLE 2. ESG LAB TEST CONFIGURATION

Two DXi7500 disk backup systems	18 TB usable, firmware 5.1.329 FC VTL emulation
Two Dell 1950 servers: Dual Xeon 3 GHz CPUs, 4 GB RAM, QLogic QLE2462 FC HBAs Driver 9.1.7.16	Windows 2003 Server, SP 2 NetBackup 6.5
One Dell 1950 server: Dual Xeon 3 GHz CPUs, 4GB RAM, QLogic QLE2462 FC HBAs, Driver 8.01.04-d8	RedHat Linux 2.6 kernel
SAN Connectivity	Brocade Silkworm 4Gb/sec FC switch
LAN Connectivity	Dell D-link DGS-3024 Ethernet switch
Exchange Data	18 full backups, 112 GB each
Server Storage	FC attached JBOD with StorNext v3.1, Adaptec arrays, EMC arrays
Tape	Quantum Scalar i500 LTO-4 drives
QUANTUM TEST CONFIGURATION	
One DXi7500 disk backup systems	72 TB usable, firmware 5.1.329 FC VTL Emulation
Two Dell 2950 servers: Dual Quad core Xeon 3 GHz CPUs, 16 GB RAM, Dual QLogic QLE2462 FC HBAs, ,Driver 9.1.7.16	Windows 2003 Server, SP 2 NetBackup 6.5
Two HP DL360G5 servers: Dual Quad Core Xeon 3 GHz CPUs, 16 GB RAM, Dual QLogic QLE2462 FC HBAs Driver 8.01.04-d8	RedHat Linux 4.5
SAN Connectivity	Brocade Silkworm 4 Gb/sec FC switch
Exchange Data	18 full backups, 112 GB each
Server Storage	4 Infotrend Arrays, 16x250GB HDDs



20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com