



Quantum Corporation

Leveraging a Unified Multi-Tiered Backup Design To Deliver a Stable, Scalable Backup Infrastructure

Abstract

Organizations face unprecedented challenges in creating a backup infrastructure that can respond to today's enterprise demands. In the face of these challenges, organizations need a backup design that unifies and stabilizes their backup environment. Quantum's portfolio of solutions solves today's backup problems while putting in place a foundation for stable and sustainable growth.

White Paper

July 2009

By

DCIG LLC
7511 Madison Street
Omaha NE 68127

402.884.9594



Leveraging a Unified Multi-Tiered Backup Design To Deliver a Stable, Scalable Backup Infrastructure

Table of Contents

| | |
|--|------|
| Executive Summary..... | Pg1 |
| Backup Redesign – The New Enterprise Initiative | Pg2 |
| Today’s Backup Problems Solved – New Opportunities Are Created..... | Pg2 |
| The New Enterprise Challenge: Managing a Multi-Tiered Backup Infrastructure..... | Pg3 |
| The Pitfalls of Creating Heterogeneous Backup Infrastructures | Pg5 |
| The Need to Create a Unified, Manageable Multi-Tier Backup Infrastructure | Pg7 |
| Quantum’s Unified Hardware and Software Suite | Pg8 |
| The Case for Quantum’s Multi-Tiered Backup Solution | Pg11 |



Leveraging a Unified Multi-Tiered Backup Design To Deliver a Stable, Scalable Backup Infrastructure

Executive Summary

The importance of protecting and recovering data is nothing new, but given the dramatic changes that organizations are experiencing in regards to data growth, server virtualization and data center consolidations, organizations are putting more data than ever at risk. These changes are contributing to a proliferation of new technologies that are dynamically reshaping data center infrastructures and prompting organizations to make similar shifts in how they backup and recover their data, and how they retain and protect it over time.

Organizations are well aware that they need to make fundamental changes to how they backup their data as backup redesign is now top of mind in a large number of Fortune 1000 organizations. The challenge becomes how to cost-effectively redesign their backup infrastructure so that it takes advantage of the latest technologies, resolves their current backup pain points, puts in a foundation that can grow with them and does so without sacrificing the stability of their day-to-day operations.

Ending the daily backup pain is often the first objective of most organizations. To do this, enterprise organizations are adding disk-based backup systems of different capacities and interfaces to their existing tape infrastructure, resulting in the creation of a multi-tiered backup infrastructure.

But what organizations may quickly find out is that if they use disk and tape libraries from different vendors or even multiple models of disk and tape libraries from the same vendor, they only create new issues longer term. Data may become siloed at certain sites; migration from disk to tape and back again becomes difficult to execute; or they discover that monitoring and managing individual disk and tape libraries becomes a task that is haphazardly addressed.

Enterprise organizations, now more than ever, need a holistic solution that delivers multi-tier backup within an integrated infrastructure. The approach needs to solve their immediate backup pain, offer solutions for both remote offices and data centers, come at price points that are sensitive to the needs of these respective locations, manage data throughout its entire life cycle, and provides centralized management.

The combination of Quantum's DXi™-Series disk libraries, Scalar® tape libraries and Vision™ management software create a unique value proposition in today's crowded backup space. Quantum's portfolio of hardware and software solutions provides organizations with the immediate backup relief that they seek regardless of where they need it.

Quantum's DXi-Series of disk-based backup systems offer deduplication and replication that make storing data on disk longer term more practical, while Quantum's tape libraries provide the needed technology and encryption for long-term data retention. By using Quantum Vision, organizations have a central console to manage these devices and the data on them regardless of where these are located within the organization.

This redesigned backup infrastructure remains cost-effective, manageable and scalable despite continued organizational growth amidst scarce resources.



Leveraging a Unified Multi-Tiered Backup Design To Deliver a Stable, Scalable Backup Infrastructure

Backup Redesign – The New Enterprise Initiative

Enterprise organizations have always recognized the importance of protecting and recovering their production data but the look and feel of the computing infrastructure of today's organizations has changed dramatically in the last decade. New and maturing technologies such as archiving, ample network bandwidth, data center consolidations, eDiscovery, explosive data growth, tiered networked storage and server virtualization are all contributing to a dynamic reshaping of the infrastructure of many organizations. But what many organizations have yet to do is make a similar shift in how they backup and recover their data.

Organizations are well aware that they need to make changes to how they backup their data as backup redesign is top of mind with a growing number of enterprise IT managers. Feedback gathered from users in Fortune 1000 organizations over the last few years by TheInfoPro has continually placed backup redesign as a leading storage initiative.

However, the redesign of backup infrastructures is being driven by more than just users wanting to take advantage of new and better technology. Other reasons include:

- ***Inefficient use of IT staff, time and dollars.*** Backup administration is the largest consumer of IT staff time and monies in organizations taking up to 20% of their time and 15% of their budget dedicated to backup and recovery.
- ***New legal search and hold requirements.*** New federal eDiscovery rules require organizations to produce electronically stored information including backup data in as quickly as 30 days. Once located, organizations may need to place a legal hold on this data and retain it until the litigation is complete.
- ***Exploding data growth.*** Data growth rates continue at 30, 40, or 50% or more annually in good economic times and bad even though storage budgets may remain fixed or even contract. These data growth rates pressure existing backup windows and recovery times.
- ***Users are coming to work with higher expectations.*** Users are finding they can now quickly backup and recover their own data at home and are bringing those same expectations into the workplace. They now expect the businesses that they work for to deliver the same swift levels of backup and recovery for their production applications.

Today's Backup Problems Solved – New Opportunities Are Created

Organizations that have already begun to redesign their backup infrastructures are placing a higher priority on using the right storage technologies that improve backup success rates, optimize storage capacity and move data between different tiers of storage. Disk backup systems, both VTLs and other disk target devices, are growing in popularity as a primary backup target as they reduce backup times while increasing the success rate of backup jobs.



Data deduplication is further accelerating the adoption of disk as a backup target. By deduplicating backup data, organizations now can store more data on disk and for longer periods of time. Deduplicating backup data also helps companies solve another problem that storing data on disk creates – data portability.

Storing backup data on disk in its native format can make it more difficult to move offsite, so organizations either need to copy the backup data off to tape or purchase additional network bandwidth to replicate it to a secondary location. While organizations may still follow a traditional disk-to-disk-to-tape strategy even using deduplicated data, the reduction in backup data that deduplication delivers makes it practical and affordable for organizations to replicate data to other sites. Organizations can use existing WAN links to replicate data without requiring more bandwidth since deduplication minimizes the amount of data being replicated, limiting the transfer to the new blocks that appear in each backup cycle.

Combining disk and deduplication also create new opportunities for organizations to protect their remote and branch offices (ROBOs). By placing these deduplicating disk systems at these sites, organizations can simultaneously improve the backup situation and set the stage to centralize data protection at the main data center. Using the optional replication feature on these disk libraries, organizations can replicate data from these remote locations back to the main data center.

Once the data is there, organizations can deploy tape libraries at the main site, the disaster recovery (DR) site, or both to keep storage costs to a minimum. DCIG expects more enterprise organizations to continue to use tape for the foreseeable future for several reasons:

- Tape provides the most affordable long term retention for archival data that is infrequently or rarely accessed
- It keeps energy expenses affordable since the power costs for disk can be as much as 290x more than tape.¹
- Tape protects against data corruption by isolating data to ensure that it is not tampered with. The capability has been recently enhanced by the cost-effective integration of encryption into standard tape drives and by the inclusion of key management capabilities from tape library vendors, including Quantum.
- Tape offers portability for offsite DR.

However, as disk-based systems and tape libraries start to proliferate within enterprises in their ROBOs, new management challenges emerge for both the main data center and DR site.

The New Enterprise Challenge: Managing a Multi-tiered Backup Infrastructure

Different capacities, models and features on disk-based systems and tape libraries create a multi-tiered backup infrastructure which creates management challenges that organizations may not foresee. Organizations will want to deploy the appropriately sized disk or tape library in each location, be it at a ROBO, data center or DR site, so different models with different performance and capacity characteristics are deployed at these sites. ROBOs will most likely install smaller disk-based systems at their sites while the main data center and DR site may both deploy fully featured, highly scalable disk-based systems and tape libraries at these sites.



Common issues that organizations will face as they look to manage individual disk-based systems and tape libraries across the enterprise include:

- Tracking which disk or tape-based system is in each physical location
- Managing and maintaining user logins and passwords
- Monitoring and managing capacity utilization and performance
- Troubleshooting backup and hardware problems
- Disk and tape-based systems each present management problems that are unique to them.

By monitoring these issues, enterprise organizations can then use this information to determine immediate and near-term problems on disk-based systems as well as anticipate what types of upgrades or changes they need to make, if any, going forward. For instance, enterprise organizations may choose to replicate deduplicated data from ROBOs back to the main data center but they do not want to disrupt production applications at either site. By monitoring the replication of the data between the sites as well as the activity on the disk-based system while backups and/or restores are occurring, organizations can determine the best and worst times to replicate data. They can then schedule replication to occur at times that do not negatively impact production applications, backups or network traffic.

Disk-based System Management Checklist

- How much disk capacity is on each disk library?
- How much of that capacity is available versus how much is used?
- What sort of deduplication ratio is the disk library achieving?
- How well is the disk library performing?
- Are there any types of loads that are causing problems?

Tape libraries present equally challenging issues. Organizations will want to minimize down time associated with tape drive and media cartridges by managing factors such as cleaning, monitoring wear and tear on tape drives, providing historical performance data, and other factors.

The tape library will also provide its own management interface that organizations use to monitor and manage it. At a most basic level, organizations will use this interface to create administrator and user accounts for those individuals who manage the tape library and the components in it.

The tape library management software may also provide important statistics on how well the components of the tape library are behaving and/or performing. On tape drives, it may gather statistics on throughput ratios, utilization and idle times as well as the number of media errors occurring on specific tape drives. On tape cartridges, the software may track the utilization on individual cartridges and the number of bad blocks/sectors on them.



Tape Library Management Checklist

- Monitor wear and tear on the tape library mechanics
- Regularly clean tape library tape drives
- Create administrator and user logins to tape library consoles
- Monitor performance statistics on tape drives
- Monitor media errors on tape cartridges and drives

Some tape libraries also are taking on more advanced responsibilities such as assuming the role of managing the encryption keys that are used to encrypt backup data. In these circumstances, organizations must ensure the encryption keys are managed and protected so when they recover encrypted data from tape locally or remotely they can provide the necessary keys to decrypt the data.

Yet what organizations will find is that managing just a few disk or tape-based systems can become very cumbersome and inefficient to monitor and manage when one must log into each one and manage it individually. Further, organizations may find that they have no efficient means to aggregate statistical data from all of these devices or, if they are using models from different vendors at their different sites, to effectively manage them either.

The Pitfalls of Creating Heterogeneous Backup Infrastructures

An organization that deploys a heterogeneous backup infrastructure can quickly discover that the initial cost savings that one may achieve from deploying disk and tape-based system from different vendors is more than offset by the costs associated with managing it.

Third Party Management Software is No Panacea

The management software on each vendor's disk and tape-based systems is very specific to each solution so it is not always conducive to being managed by third party software for any number of reasons. For example, there is a large number of disk and tape-based systems that third party management software products need to account for. Therefore it is not a given that a third party software package that promises to manage heterogeneous backup infrastructures can necessarily deliver on that claim.

Challenges with 3rd Party Management Software

- Large numbers of 3rd party devices to manage
- Support often lags release
- Need to justify value of 3rd party software
- Quality of relationship between hardware and software providers



Even in circumstances where a third party software product does manage a large number of devices, it will need access to the APIs of the management software on those disk or tape-based systems in order to manage them and the correct version of that software. This is easier said than done as new APIs emerge over time. How well and how soon third party management software can take advantage of these new features will depend on how quickly the software provider integrates that functionality into their software and, to a certain degree, the relationship that exists between it and the providing hardware vendor since cooperation will be required to deliver on those functions.

Storage Standards are Struggling at Best

Storage standards are also sometimes pointed to as a mechanism by which third party software can manage disk and tape-based systems from different hardware vendors. However this initiative has lost much of its momentum over the last few years and even the progress made to date is dubious at best. The demand from both organizations and vendors to deliver an industry standard interface for disk and tape-based systems' management is mixed at best.

Challenges with Storage Standards

- No current industry push for wide scale adoption of standards
- Generally only provides basic levels of reporting
- Minimal or no recognition of advanced features on disk or tape-based systems

Standards that do exist only provide more basic levels of storage reporting and management. For instance, standards may allow third party software to report on storage utilization on the disk-based system but do not recognize utilization on specific partitions within the disk-based system or it may not even recognize that partitions exist and so it cannot manage them. These base-line levels of reporting and management therefore are inadequate to deliver the sort of functionality that enterprise organizations may want and need.

Advanced Software Features do not Interoperate

The problems of heterogeneous backup infrastructures only become more aggravated when advanced features such as replication and encryption are used. These can create interoperability issues that mitigate some of the anticipated cost savings that organizations may expect when they put them in place. For instance, a software feature such as replication is incompatible between disk-based systems from different vendors. This makes it impossible to replicate data between these different vendor's respective disk-based systems.

This same principle also holds true for encryption on tape libraries. Data encrypted by one vendor's tape library may be inaccessible and unreadable on another vendor's tape library since the keys used to encrypt the data on one system can not be shared with the other tape library.



So to effectively use encryption or replication on an enterprise scale and realize its full benefits, enterprise organizations must use the same vendor's software at all of an organization's sites. The difficulty in achieving these objectives is identifying a vendor that provides the breadth of products that meet the specific but very different needs of both ROBOs and data centers.

Enterprise Demands

Data centers require sophisticated disk and tape-based solutions that can meet their demands for high performance, high capacity and advanced management functions. These solutions require dedicated IT staff that has the background, expertise and time to configure and manage these devices in accordance with the requirements of the business and its applications.

ROBOs require disk and tape-based systems that meet a very different set of demands. ROBOs may have little or no expertise or dedicated IT staff on hand so they need solutions that require little or no management and ideally can be managed by IT staff in the main data center. Also, since their backup and recovery needs are less demanding than those found in data centers, they want solutions that are cheaper and easier to manage but can still provide some of the advanced features such as replication and encryption. Ideally these solutions in the remote offices can be managed in the same manner as the larger, more robust disk and tape-based solutions that are found in the main data center are managed and can interact with them.

The Need to Create a Unified, Manageable Multi-Tier Backup Infrastructure

This dichotomy that exists between the product and IT staffing needs of ROBOs and data centers has created a dilemma for enterprise organizations in identifying and implementing a manageable end-to-end backup infrastructure. Essentially, enterprise organizations want:

- ***Centrally managed disk and tape-based systems.*** Organizations want to reduce the amount of IT staff time that is spent managing backup and recovery from current levels to ideally 5% or less. To achieve this objective, they must put in solutions that they can centrally manage anywhere in the enterprise that improves the current backup and recovery situation.
- ***Interoperability between disk and tape libraries.*** Disk-based systems are assuming a larger role in more enterprise organizations as a primary backup target and for short term data retention (30 – 90 days). That does not diminish the role that tape libraries will continue to play in many enterprise organizations, especially in their data centers. However these organizations will want a single interface to monitor and manage both disk and tape-based systems and then move data between them. This may include moving data from disk-based systems in remote offices to their data centers and then back and forth between disk-based systems and tape libraries in the data center.
- ***Appropriately sized and priced disk and tape-based systems for ROBOs and data centers.*** Enterprise organizations obviously want to purchase the right size disk-based system or tape library for the environment into which they are to be deployed. They will want smaller, simpler and less expensive disk and tape-based solutions for ROBOs while they will be willing to pay more for the more robust solutions that would be installed in data centers.



- **Advanced features that work across all disk and tape-based solutions.** Replication is one software feature that organizations will want to utilize in ROBOs and data centers as DCIG finds that up to 50% of users deploy replication as part of their disk-based solution. Yet centrally managing the replication of data is only possible if the software works across all of these platforms.
- **Centralized reporting and management software.** Organizations will want to track in each location how well the disk or tape-based system is performing, how much capacity is used and forecast growth rates as well as make changes to it as needed. To do this on a wide scale across all disk and tape-based systems across the enterprise is only possible if organizations have access to software that can centrally monitor, manage and report on these different devices.

Quantum's Unified Hardware and Software Suite

This need for an enterprise, multi-tiered backup infrastructure is the void that Quantum's unified suite of hardware and software products seeks to fill. Quantum offers disk-based systems for both ROBOs and data center and DR sites and tape automation solutions for long term data retention.

Quantum's DXi2500-D disk-based backup appliance is intended for use in ROBOs and comes with the hardware and software features that one might expect in a disk-based system intended for these cost sensitive and less demanding environments. The deduplication, management and replication software features available on the DXi2500-D are compatible with the same features found on Quantum's larger DXi7500 disk-based backup system. While they do not have all of the advanced options that the DXi7500 offers, organizations can replicate data from its DXi2500-D to the DXi7500 and manage them in a comparable fashion.

Quantum differentiates itself from some other disk-based system providers in that it also offers its Scalar line of tape libraries that provides an integrated method to move data from disk to tape. Organizations may want to take advantage of this to move their aging backup data from disk to tape for long term data retention or offsite data storage. In so doing, they can obtain both their disk and tape-based systems from Quantum so they have one point of contact.

However Quantum's robust portfolio of disk and tape-based solutions demand organizations consider Quantum as their sole provider for disk and tape-based solutions in order to achieve enterprise data protection. Specific features that set Quantum's disk and tape-based systems apart from competitors include:

- Quantum's DXi7500 interacts with its Scalar tape libraries so it provides organizations with an integrated path from disk to tape. Organizations can leverage the DXi7500 in the main data center, to copy the data from disk to the Scalar tape library. The DXi7500's path-to-tape feature integrates with existing enterprise backup software products such as Symantec NetBackup and EMC Networker so that it can track in its catalog what data is located where – whether it is on disk, tape or both.
- The Quantum Scalar Key Manager Appliance is an optional feature that is available on Quantum's Scalar i2000 and i500 tape libraries. By leveraging this feature, organizations have a built-in strategy to encrypt data stored on their LTO-4 tapes and then manage the encryption keys longer term.

DCIG

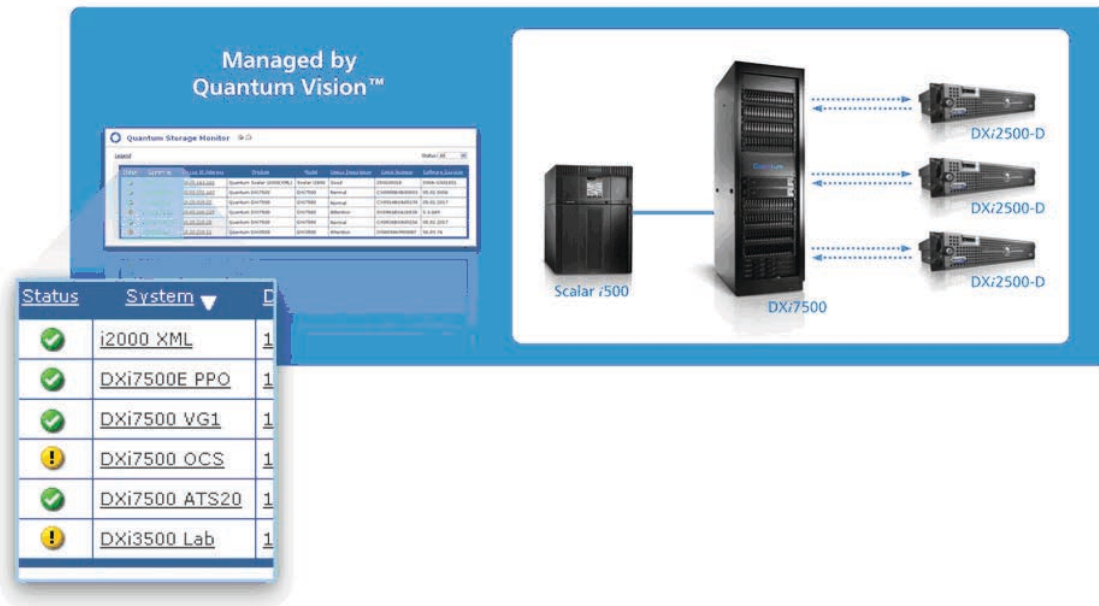


Figure #1 – Quantum Vision software enables organizations to consolidate the monitoring and management of their disk and tape systems across different sites.

It is Quantum’s Vision software that brings the entire story of disk and tape together so organizations can deliver unified management across their backup infrastructure. Quantum Vision can be used to monitor, manage and report on any DXi-Series disk backup systems or Scalar tape library regardless of where it is located assuming that a network connection is available to it.

The Quantum Vision software eliminates the need for administrators to log into individual disk and tape-based systems to perform specific management tasks, instead providing a central console from which all of these devices are managed. It also aggregates statistics such as deduplication ratios, disk/tape utilization and trending information so organizations can get a complete picture of their disk and tape environment.

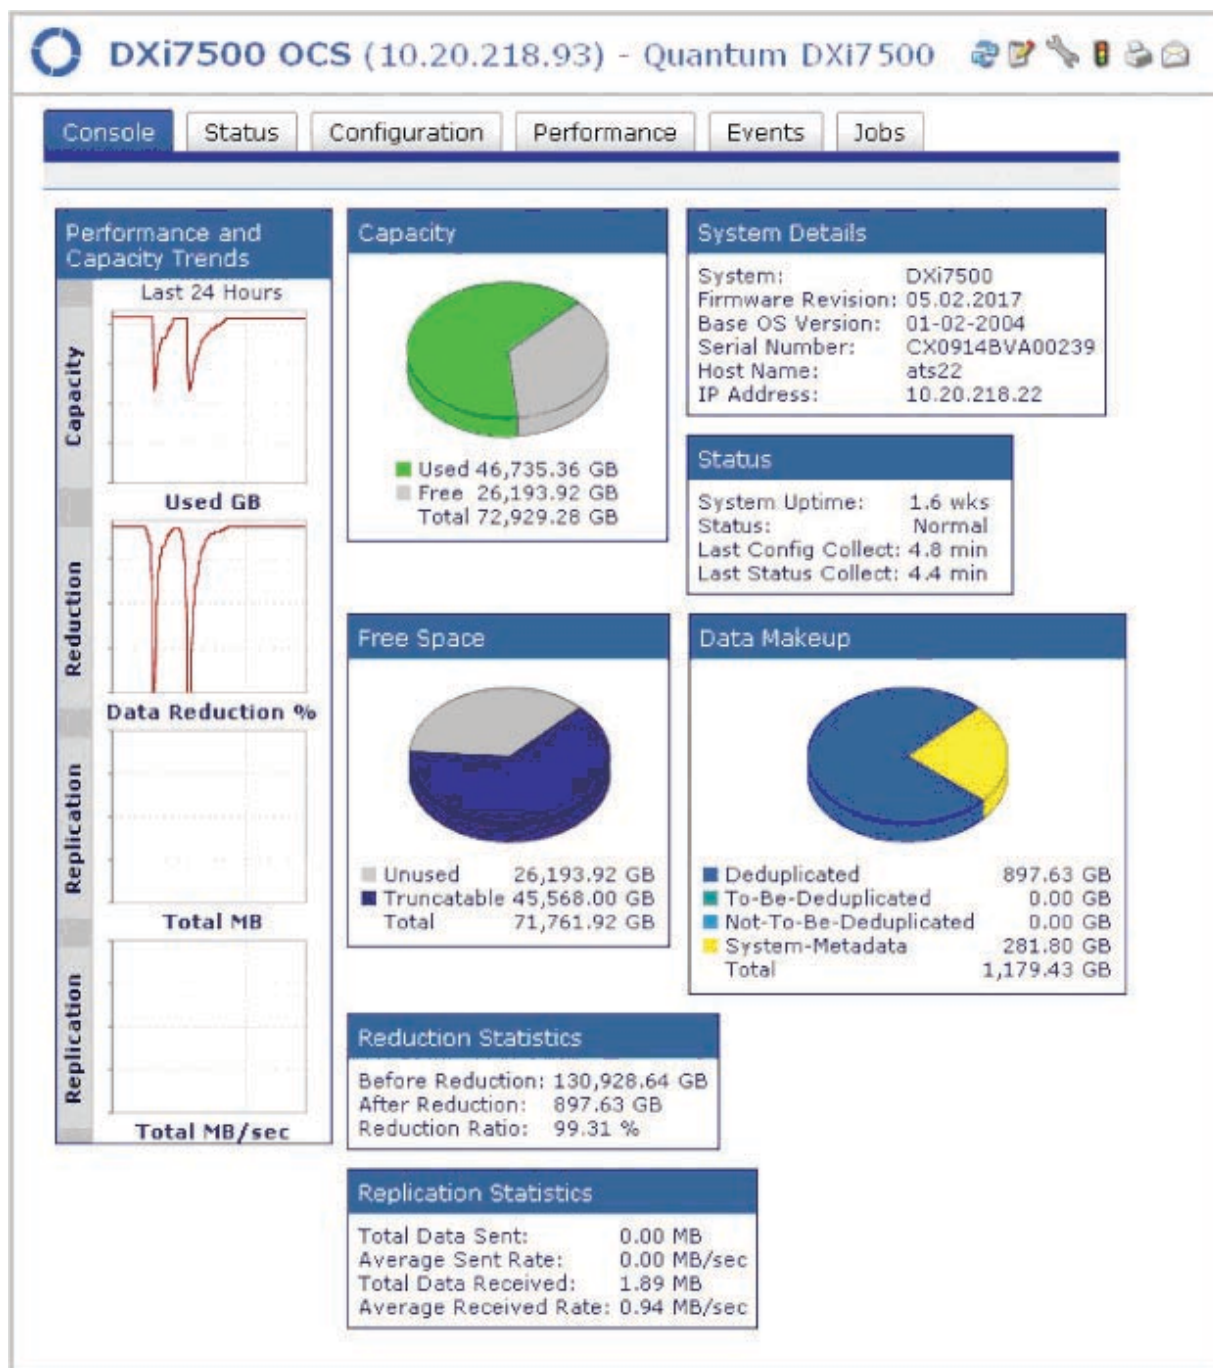


Figure #2 – Quantum Vision enables organizations to analyze critical data across multiple units and over time. The above screen shot highlights capacity utilization, deduplication ratio, system status and configuration.



Because it is in communication with these devices, Quantum Vision can also proactively send out notifications should problems occur on any of these disk or tape solutions. For instance, on disk-based systems it monitors and sends out notifications when maximum capacities are reached or policy-defined thresholds are crossed. On tape libraries, it looks for excessive error rates on tape drives and tape cartridges (the leading causes of failure). In both cases, administrators can be notified before any hardware failure or out-of-space condition occurs so they can take action.

The Case for Quantum's Multi-Tiered Backup Solution

The combination of Quantum's DXi-Series disk-based backup systems, Scalar tape libraries and Vision software creates a unique value proposition in today's crowded backup space. While enterprise organizations may initially focus on deploying disk-based systems into their data center, ROBOs and DR sites to address immediate backup challenges, they can inadvertently create a heterogeneous backup infrastructure that is difficult to manage and navigate. In so doing, organizations create a larger management problem that fails to take into account how they are going to manage their disk and tape-based systems and the backup data stored on them.

A better approach is for organizations to implement a solution that addresses their immediate backup pain while also putting a foundation in place that avoids these future management challenges. Quantum's portfolio of hardware and software solutions provides organizations with the immediate backup relief that they seek regardless of where it is needed.

Since Quantum uses the same software across its DXi-Series disk-based systems and Scalar tape libraries, organizations can take advantage of features such as deduplication, replication and encryption while still creating a central console to manage these devices and the data on them regardless of where they are located in the organization.

Quantum's solutions keep tape a part of the organizational backup strategy. By enabling its disk-based systems to interact with certain tape library models and configuring its Vision software to manage both its DXi-Series disk systems and Scalar tape libraries, Quantum delivers to organizations an enterprise wide multi-tiered backup infrastructure. This redesigned backup infrastructure remains cost-effective, manageable and scalable despite continued organizational growth amidst scarce resources.

¹Disk and Tape Square Off Again – Tape Remains King of the Hill with LTO-4; <http://www.clipper.com/research/TCG2008009.pdf>; David Reine and Mike Kahn, The Clipper Group; February 13, 2008.

NOTICE: The information, product recommendations and opinions made by DCIG LLC are based upon public information and from sources that DCIG LLC believes to be accurate and reliable. However since market conditions change, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. DCIG LLC assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise) caused by one's use or reliance of this information or the recommendations presented or for any inadvertent errors which this document may contain. Any questions please call DCIG LLC at 402.884.9594.