

Energy and Space Efficient Storage: Multi-tier Strategies for Protecting and Retaining Data

NOTICE

This White Paper may contain proprietary information protected by copyright. Information in this White Paper is subject to change without notice and does not represent a commitment on the part of Quantum. Although using sources deemed to be reliable, Quantum assumes no liability for any inaccuracies that may be contained in this White Paper. Quantum makes no commitment to update or keep current the information in this White Paper, and reserves the right to make changes to or discontinue this White Paper and/or products without notice. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any person other than the purchaser's personal use, without the express written permission of Quantum.

CONTENTS

Power, Cooling and Space Concerns in the Data Center	3
Impact of Trend on Disk Backup	3
Data Deduplication Is Key Enabler	3
Meeting Requirements with Multiple Tiers of Storage	4
Summary	5
Quantum Solutions for Multi-Tier Data Protection and Retention	6

Data center power, cooling, and space requirements are becoming key IT challenges for organizations of all sizes. At the same time, IT managers are faced with ever-increasing demands for data protection, including better restore performance and longer retention times. Meeting these challenges requires a multi-tier storage architecture along with new technologies like data deduplication.

Power, Cooling, and Space Concerns in the Data center

Power, cooling, and space requirements are key issues faced by IT managers. Demand-side contributors to this situation include decreased hardware acquisition costs, sharp increases in server density, and an exponential rise in the volume of data being stored and managed.

The Environmental Protection Agency has predicted that data center power usage will double over the next five years. The EPA released a paper by Dr. Jonathan G. Koomey, professor at Stanford University, stating that 1.2% of all power purchased in the United States was consumed by data centers in order to operate and cool equipment¹. He also warned that this phenomenon is on the rise and could impact data centers worldwide if we don't focus our attention toward green computing.

On the supply side, electric utilities already struggle to keep up with demand. Brown-out restrictions can be commonplace during peak consumption periods, and many IT directors are finding that they cannot increase the amount of power they source from their local power grid. IT analysts believe that by the end of the decade the world's data centers will have run out of power. Yet, building new, more efficient facilities can be cost prohibitive due to construction and cooling costs.

While data centers remain an essential part of most businesses, there will come a time very soon when growth may need to be limited because of power costs and the growing awareness at governmental, corporate, and personal levels in regards to the importance of reducing our environmental impact.

Impact of Trend on Disk Backup

The growing restrictions on power, cooling, and use of data center footprint can be seen as coming into conflict with another major IT trend namely a steady increase in the percentage of disk in the data protection mix. Using disk as part of a backup strategy is growing rapidly as IT departments take advantage of the faster backup and restore performance that disk delivers. In turn, disk provides higher service levels to their customers and reduces the use of removable media. But disk storage is fast becoming one of the primary power and cooling culprits in the data center, and when disk is used for backup, the requirements can increase dramatically.

The solution is not to stop data protection progress, but rather to take advantage of various storage technologies by combining different kinds of media for an optimal balance of performance and cost, including the ongoing energy consumption factor. One of the new technologies that is having a major impact in this area is data deduplication.

Data Deduplication Is A Key Enabler

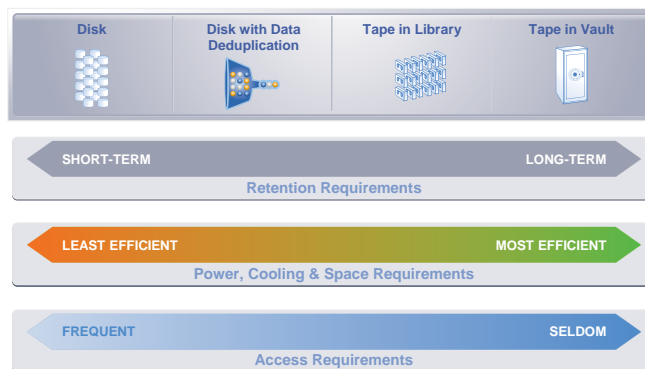
Data deduplication technology is being vaunted by industry analysts as the next big IT advance, with the Taneja Group predicting 108% compound growth for this market sector through 2010. There are several reasons for this prediction. Data deduplication reduces the amount of disk required to protect

a given amount of primary data by 90% or more by detecting and eliminating redundant blocks within files and, in some products, between different files and file types. It allows users to exploit the performance characteristics of disk for backup without incurring the costs of conventional disk systems—costs which include both capital expenses and the operating expenses that include power and cooling.

Disaster-recovery protection also benefits from some vendors’ implementations of data deduplication. Using the same technology that identifies duplicate segments inside datasets, data deduplication systems also reduce the bandwidth needed to transmit backup sets over a network. With data deduplication, once systems are synchronized, whole backup sets can be replicated while only changed blocks are actually moved. If a new backup is only 5% different from a previous one at a block level, then we can reduce the bandwidth needed for transmission by up to 95%. Replication allows multiple copies of backup data to be maintained on disk in different sites using WAN links. Its use can simplify backup and reduce costs by decreasing an IT department’s need to buy, manage, and move tapes. At the same time, it provides improved disaster recovery protection.

Meeting Requirements with Multiple Tiers of Storage

The optimal approach to data protection and retention is to match the various business requirements with the right storage technology. IT managers must consider the business requirements for data access and retention, and balance those against the requirements for power, space, and cooling. In most cases, this will result in a multi-tier storage architecture that includes a mix of disk, data deduplication, and tape technologies.



Several factors must be considered prior to implementing a backup, recovery and archive plan.

Datasets with high access requirements may need to use conventional disk systems without data deduplication. However, they are only practical for a subset of data where the recovery time objective overrides other considerations such as long-term retention and total cost. Datasets with normal access requirements can use disk with data deduplication technology both as first site for backup data and for medium term retention. Data deduplication technology reduces space, power, and cooling requirements enough to make it practical to use disk as a retention medium for weeks or months without breaking the operational budget. However, as data needs to be retained for multiple quarters and years and the demand for fast recovery diminishes, the most effective retention medium for most users is tape. Tape cartridges in a tape library consume power at lower rates than any disk system, and tape cartridges stored in a vault consume the least of all, as well as providing the lowest aggregated storage cost per GB. The following chart compares these various types of storage in terms of power, space, and cooling.

	Disk Array ¹	Tape Library ²	Saving	Disk with data deduplication ³	Saving
Space	0.01m ² /TB	0.00215m ² /TB	78%	0.0003m ² /TB	97%
Power	60 Watts/TB	.88 Watts/TB	99%	3.47 Watts/TB	94%
Cooling	204 BTU/TB/H	3.0 BTU/TB/H	99%	12.76 BTU/TB/H	94%

1. EMC CLARiiON CX20 with 120 750GB drives
2. Quantum Scalar i2000 with 300 slots, LTO-4, 2:1 compression
3. Quantum DXi3500, 10.8TB, 20:1 data deduplication

Summary

IT managers are faced with growing concerns around data center power, cooling, and space requirements. Costs are on the rise and supply is becoming increasingly scarce. At the same time, IT managers must deal with ever-growing demands for protecting and retaining valuable data. Meeting these various demands requires a multi-tier storage architecture with the appropriate mix of disk, tape, and enabling technologies like data deduplication.

Quantum offers products that fit into all of the different categories, and offer centralized management to make it easy to discover, manage, and monitor multiple tiers of storage from a common console. Quantum takes a solution-based approach to these challenges, offering a rich balance of scalable hardware and software systems plus a variety of value-added services, to meet data protection and retention requirements.

Quantum Solutions for Multi-Tier Data Protection and Retention

Quantum offers Disk-to-Disk-to-Tape (D2D2T) backup solutions: disk-based backup that includes patented data deduplication and replication (DXi™-Series) as well as a full range of tape libraries (Scalar® Series). For additional security, the tape libraries now offer encryption of data to be stored off-site and include key management. The company has recently made it easy to combine all of these products together in a single system by introducing a common management tool—Quantum Vision™—that discovers, manages and monitors all of these systems from a common console. In addition, all the systems can also be covered by a single service and support organization, which makes a common software system available for remote monitoring and diagnostics.

Even at an individual solution level, Quantum is offering an integrated approach to multi-technology systems. The DXi7500, for example, is a high-availability enterprise disk backup system that combines multiple technologies at a solution level. It can operate as a disk backup system using a conventional virtual tape library interface to provide the very highest throughput, as well as a data deduplication enabled disk backup and remote replication system. The DXi7500 also features a direct tape creation capability that allows users to automatically migrate the backup data stored on disk to tape for longer term retention—and the movement takes place in the background without impacting the user's media server or backup SAN. The DXi7500's tape creation works with any backup software, and it has direct support for all major backup software providers that have built in direct-tape creation functionality.

When you add encryption during replication (a core DXi feature), encryption on tapes for offsite storage, and common management across disk, tape, midrange, and Enterprise systems, end users have a highly integrated solutions set that makes it easy to combine the right balance of performance and cost over the life of a user's backup and retention.

Moving outside the realm of backup to archive, Quantum also provides automated migration of data across different tiers of storage for primary data with its StorNext® data management software. StorNext can create multi-tier archives for unstructured data pools employing tier 1, and secondary disk tape and deduplicated disk layers. It can even utilize existing 3rd party NAS storage as a layer within the archive. Files can be stored on lower cost storage over time. Data management is automatic and data is even protected by default, negating the need for separate backup solutions and removing the backup window. StorNext helps customers build an infrastructure for consolidating resources so that workflow operations run faster and maintaining business assets costs less.

1. Jonathan G. Koomey, Ph.D,
"Estimating Total Power Consumption by Servers in the U.S. and the World," February, 15, 2007.

Quantum is a member of the SNIA's Green Storage Initiative



For contact and product information,
visit quantum.com or call 800-677-6268

Quantum®

Backup. Recovery. Archive. It's What We Do.

©2009 Quantum Corporation. All rights reserved. Quantum, the Quantum logo, and all other logos are registered trademarks of Quantum Corporation or of their respective owners. Protected by Pending and Issued U.S. and Foreign Patents, including U.S. Patent No. 5,990,810.

About Quantum

Quantum Corp. (NYSE:QTM) is the leading global storage company specializing in backup, recovery and archive. Combining focused expertise, customer-driven innovation, and platform independence, Quantum provides a comprehensive range of disk, tape, media and software solutions supported by a world-class sales and service organization. This includes the DXi™-Series, the first disk backup solutions to extend the power of data deduplication and replication across the distributed enterprise. As a long-standing and trusted partner, the company works closely with a broad network of resellers, OEMs and other suppliers to meet customers' evolving data protection needs.