

Deduplication in the Enterprise Data Center:

Assessing the Impact and Benefits

Table of Contents:

Quantum Interview

Page 5
Key Elements of a Successful Storage Strategy

Page 12
About Quantum

Interview



Janae Stow Lee is the Senior Vice President of Disk and Software at Quantum Corporation. She has over 30 years of experience in the storage industry, including ten years working with a variety of data reduction and file system software and hardware. Prior to Quantum, she was CEO at TimeSpring Software (now Doubletake Software); VP of Product, Marketing and Business Development at Avamar (now EMC); and a senior sales and marketing executive at both Legato Systems (EMC) and IBM Storage.

In a recent interview she responded to several questions about the many benefits IT and Storage Managers are finding as they have deployed deduplication in their data centers.

What kind of impact have you seen deduplication having on the way that IT departments protect their data?

Even though I've been doing it for several years, I still find myself being impressed when I talk to our DXi customers and they tell me how much their general backup processes have improved since they added deduplication appliances. It's natural to expect to see first stage backup improve and it really does – it's common to have customers tell us they have doubled or tripled their backup speeds just by adding the appliance. I think it would also be easy to expect that restores are more efficient, and that is born out, too – the most common report is that restores used to take a few hours and now they happen in a few minutes. What we were a little surprised at was the scale of reduction in backup management time – people are reporting a nearly two-thirds reduction overall in the total time that they spend managing their backup and recovery process. That is a massive gain – for most users, it means that significant IT resources are freed up for other jobs.

Cost reduction is another huge benefit – it applies to customers who just use the appliances to do backup and restore, but it is multiplied even more when users put in replication to provide their first stage DR. One of our customers, Barclaycard, just won a Best Practice Award at the Fall 2010 SNW and a major reason was cost savings. They put in a DXi system that uses disk for backup and restore, uses replication of backup data for lights-out, day-to-day DR protection, and writes data to tape for long term retention. Their summary said that they reduced media costs by 77%, saw an 80 to 85% reduction in media management hours, cut their off-site storage provider costs by 68%, and eliminated an annual six-figure charge from a third-party handling company – total annual savings were \$300,000. At the same time, they improved their protection – they have DR protection every day – and in the question and answer period after the presentation, they estimated the financial payback period at 13 months. It's outcomes like these that are helping accelerate a kind of new wave of deduplication adoptions – one based on proven outcomes more than just on cool technology factors.

It seems as if everybody is talking about deduplication today, especially backup software suppliers. How much of a shift is going on to move deduplication into backup applications?

Whenever there's a new technology with real value, people experiment by trying slightly different approaches and by locating at different places. We saw that a few years ago with compression – there are many different places you can compress data and people have tried different approaches to see what the tradeoffs are. The same thing is happening with deduplication. The reality of deduplication is that to do it effectively requires lots of processing power and doing it at high speed, as Quantum has figured out over the last few years, this means being smart about how you integrate software operations with hardware. That's why the early success in deduplication has come from people who put it on purpose-built appliances.

When you move the technology into pure software applications, you end up carrying it out on a wide variety of general purpose server platforms and storage system. That makes the results a little less predictable, and at times the impact on backup performance can be very negative. It makes sense when you think about it. All the deduplication appliances include both deduplication and

compression – in the past, just compression data in backup software has slowed down backups. When you add a completely new, and even heavier process, the impact can be quite severe unless people have new-generation hardware, dedicated servers and storage, and integrate it very well. We've seen that impact even for the backup apps that use less powerful forms of deduplication to minimize loading. Is there a place where deduplication makes sense in the backup application? Absolutely! But there are also many, many sites where the cost to deduplication in the application – and that cost includes performance, integration overhead, and deduplication effectiveness, as well as dollars – makes appliances the better option. The current wave of deduplication adopters that we're seeing includes those who tried the software approach and are switching now to appliances.

There is also quite a bit of discussion about deduplication that divides activities between different elements – hosts and servers, backup servers and appliances, that sort of thing. What's your sense of that general approach now and in the future?

Customers continue to be challenged to manage data growth within constrained budgets. Data reduction technologies, including deduplication, are therefore valuable across the information lifecycle. This has driven – and will drive the industry to deploy these technologies in every choice of architecture. This is generally good news for customers – choice is a good thing.

But customers, as always, need to be concerned about balancing the value of a technology against the cost to deploy and manage it.

First, the value of deduplication is driven by how much duplicate data there is to extract, and what the value is of extracting it. If there's no gold, you shouldn't go mining! Deduplication is so widely adopted in backup because backup natively creates a lot of duplicate data; and because backup resources – the storage network, production server channels and backup windows represent very scarce assets. Secondly, value is based on how effective a particular deduplication technology is at extracting the

Deduplication in the Enterprise Data Center: is published by Quantum. Editorial supplied by Quantum is independent of Gartner analysis. All Gartner research is © 2010 by Gartner, Inc. and/or its Affiliates. All rights reserved. All Gartner materials are used with Gartner's permission and in no way does the use or publication of Gartner research indicate Gartner's endorsement of Quantum's products and/or strategies. Reproduction and distribution of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

extra data – how good is the mining process? It's been proven many times that over the range of customer backup data, a variable-length sub-file deduplication approach, such as Quantum's, is the most effective – it generates the highest deduplication ratios across a range of data types, resulting in the highest data efficiency. Finally, customers need to consider how any vendor's total solution integrates into their environment. For example, for customers who intend to continue to use tape as a third tier of storage, solutions that leverage tape have inherently higher value than those which don't. Solutions that leverage the best features of the customer's backup software to make the environment easy to manage (such as storage lifecycle policies in OST) also have extra value.

Now let's look at the cost side of the equation. Cost isn't just about the acquisition price or even the service pricing. When you deploy a new technology, you also need to consider costs like disruption to the existing environment to deploy it – if you have to upgrade a large portion of your existing environment to use it, that's a cost. Running the deduplication process is not free – it does take processors and memory to run it. This means either buying a solution optimized to run it like a target appliance or figuring out how the added impact will change your existing backup architecture. This change introduces new cost – not only in hardware, software and labor, but in opportunity. What else could you be doing with those storage engineers? The new appliances we see coming from formerly software-only backup companies is a real symptom of the cost and complexity that comes from adding deduplication in backup software. And increased management complexity also represents higher cost. If deploying deduplication means you need to implement multiple approaches to backup, there is administration complexity to consider. Solutions that integrate seamlessly and within existing administration skills do have lower cost.

Cost can also come in the form of latency. As I said before, deduplication is not free – but in addition to resources, the process also takes a certain amount of time. Putting a deduplication equation into every storage transaction means extending the transaction length – and the time to a response. For backup, where throughput is much more critical than latency, this is less important. But the need for zero-added latency is the reason why deduplication has – to date - not been adopted broadly in primary.

There's another tradeoff that deserves mention here – turning data reduction technology on in primary can reduce the value of deduplication 'downstream' in archive and backup, where it may have less risk and more value. The easiest example of this is compression – if you compress data sets in primary to get a 2:1 reduction, you're

less likely to get as much benefit from deduplication and compression in backup, where reduction ratios can be much higher. Or, you could be stuck 'rehydrating' data in between tiers to achieve the maximum reduction. This, again, represents cost.

Which brings me to the future; all these technology tradeoffs are not issues customers really should have to care about. We vendors need to be smarter about how we offer deduplication, to natively balance technology tradeoffs like resource requirements against value, to give customers the optimum solution across tiers and across time, integrated seamlessly into the environment. This is the highest value. And as the value leader in data deduplication, Quantum will provide such a solution

How much value do you see deduplication providing for archive applications and what will need to happen to get there?

Data reduction is actually already pretty widely used in archive – in the form of single instance store of documents. But more sophisticated forms – like deduplication – can have value – certainly not the same level as in backup – but still value, nevertheless. How much varies widely by the source of the data and its original use case. If the original data was highly transactional, the value may be small. But if the original data was productivity data, the customer may achieve reductions of up to 70%. The source of this benefit is the human need to copy. Think about it – how many times have you forwarded an existing email or document with a few modifications? Every time we edit and forward, we create duplicate data – data that isn't removed with single instancing.

Cloud computing is becoming a hot topic. How does deduplication relate to the cloud? Does it make it more, or less important?

Deduplication could be THE key enabler to cloud storage – or, looked at in reverse, cloud storage applications could be the so called 'killer apps' for primary and backup deduplication. Unlike in a local environment, the cloud introduces its own latency issues – and as we've discussed, latency is the enemy of response time. Rather than introducing latency, deduplication can reduce latency in the cloud by reducing the amount of data that needs to be transmitted (or by identifying that certain data does not need to be transmitted at all). As with local architectures, it will be natural to introduce deduplication for backup first, but dedupe will be a valuable tool for archive and primary in the cloud as well. Deduplication and the cloud are symbiotic – like peanut butter and jelly, you can have each one alone but together they are just much better.

What's new at Quantum in deduplication? I know the company has been announcing new products – what does the line look like now?

Quantum has had a very busy year in the data deduplication arena. Within a 12-month period we have refreshed our entire DXi deduplication product line. At the end of 2009, we introduced a whole new family of midrange appliances optimized for Ethernet environments – NAS and Symantec OST backup. We followed that in the Spring of 2010 with a new set of appliances for SMB and remote office protection. This summer we introduced a new midrange deduplication appliance with a VTL interface designed for customers with Fibre Channel SANs. And in October, we announced a new Enterprise product – rounding out the transformation.

These products are all on the latest generation hardware platforms, and most take advantage of a new way of using different storage elements for different data and operations – what that gives us is extremely powerful deduplication at much faster rates of backup and restore than earlier generation products. And we've put emphasis during the development cycle on keeping the systems simple and getting maximum value. Just one example – we got rid of the complex software licensing processing that most use in this industry and we include all the software features for all these appliances in the base price. So if users decide they want to start replicating next year, or start using OST, or use our direct tape creation feature – they already

have licenses for all that software. When we look across the product line, we have performance that is equal to or better than competitors, and costs that are consistently lower. And now that we have the line on new hardware platforms, it opens the way for the next round of software enhancements.

How about the future? What direction is Quantum taking and what can we expect to see over the next year?

You'll see us focused on the data challenges our customers care about – managing data growth; living inside budgets and helping storage pros be more responsive to their users. Just as Quantum isn't just about tape – we aren't just about deduplication either. As you know, we own the world's most scalable file management and archive system for streaming data, StorNext. We own a number of data movement technologies. And we have a bunch of intellectual property around efficient storage and data management. You can expect us to take these assets and – in combination with our partners - build them into a broader set of solutions that continue to provide benefit to customers – in an integrated but open way. We will continue to be the value leader, where value is defined as total benefit of ownership to the customer.

Source: Quantum

Key Elements of a Successful Storage Strategy

The development of a comprehensive storage strategy is central to the ability of infrastructure and operations (I&O) to provide effective storage services to internal clients and reflect organizational business strategies. This research advises I&O managers to include key elements to develop and execute a successful strategy.

Key Findings

- Storage services are strategic, while suppliers' products and technologies are not. Storage teams need to understand the differences between strategic and tactical solutions, and address them accordingly.
- Over time, some technologies, products and vendor relationships may become outmoded and out of line with organizational goals.
- The creation and ongoing maintenance of a strategic plan will take a significant amount of time and resources. It may involve legal and administrative expenses.
- Assume that creating an initial strategic storage plan will take weeks or months.

Recommendations

- Align the storage strategy with the organization's goals, not with IT's goals. Invite representation from the other business units.
- Once a storage strategy has been defined with the key decision makers, solicit input from all IT groups on the selection and implementation of tactical solutions.
- Create a data management strategy that reflects the value of your organization's data, and that defines and sets archiving and data retention policies.
- Ensure that strategic plans and directions are separated from tactical maneuvers, products and procedures, therefore avoiding unnecessary dependencies on certain vendors or products.

ANALYSIS

Storage management has become an increasingly complex and critical competency. Organizations need to deal with soaring amounts of data on a tightening budget, while legal and corporate governance issues can complicate the processes and infrastructure. Meanwhile, new technologies and varied computing styles create a myriad of choices – and many opportunities for missteps. A storage strategy defines the direction as determined by the business drivers.

There are many components, and some of the high-level directions will be balanced against costs, flexibility, performance, availability, and legal or statutory requirements. Some organizations may have cost as the dominant driver, while others may have flexibility and the ability to rapidly adapt to changing business or organizational requirements as a chief priority. While storage strategies will differ among organizations, they will predominantly align with the business drivers, not technological or supplier-determined factors, which are set at the tactical level.

1.0 Establish a Dedicated Storage Team and Utility

Organizations should put into place a dedicated storage team that will own the development, deployment and maintenance of the storage strategy. This team, which must have senior IT management support should include the IT director/manager, the storage manager, the CTO and/or the chief storage architect. Its purpose is to bring about efficiency and promote the organization's goals. By managing storage issues, the team will create confidence in the rest of the IT group that storage is being well-handled and will eliminate confusion about who should make decisions about storage.

The mind-set of the storage team should be the delivery of storage as a set of services to the rest of the organization – what Gartner calls a "storage utility." Organizationally, this

team is a peer to other specialty teams, such as networking, servers, security, asset management and database administrators, all of which provide services to each other. These utilities need formal definitions of services and service levels, and formal processes for requests, interactions and chargebacks. However, some tasks will require coordination across organizational bounds.

Action Items:

- Organize a dedicated storage team comprised of a mix of technical, financial and leadership skills. Plan for extensive training over time.
- Form the hierarchy carefully, and formalize arms-length, cooperating processes.

2.0 Seek Participation From Key Groups

The storage strategy should invite input from several groups:

- **Senior management.** Support from this group is critical to success, because management will fund and champion the effort. It is vital that senior management recognizes the value that storage brings, and that the strategic storage plan be a central part of business strategic planning.
- **IT organization.** A storage strategy needs to fully integrate with a wider, business-based IT strategy and, ultimately, help the organization achieve its goals. It is important to ensure that sufficient effort is placed on building a cross-functional IT team that meets regularly to brief, update, review and make decisions. Ensure that there is no duplication of effort, such as surveying business unit needs. In the initial phases of developing the strategy, this team should take a strategic view, not a

tactical one. Gartner advises that strategic and tactical meetings with the IT team be kept separate, and that strategic meetings include representation from key business units. All IT operations groups, including the storage team, should ensure that strategies are complementary. For example, most IT operational areas aim to reduce costs. However, some strategies will be specific to the operational group, such as data integrity maintenance and data protection for the storage group.

- **Other parts of the business.** A strategic plan should begin with meetings with executives the highest level of the company and cascade to the business unit leaders to include their input. Even though storage may not be a large part of the IT budget and a small part of the organization's overall spending, it is a key and growing component. The data stored in the storage infrastructure is often the lifeblood of an organization and, therefore, a significant and critical resource.

Collect and document feedback and information from all stakeholders. Be sure to glean answers to these questions:

1. What is the overall strategic direction of the company?
2. What business metrics are most important – growth, agility, profitability, cash flow, cost containment or compliance?
3. How competitive are the business units relative to the market?
4. How are other competitors innovating with technology?
5. How are business unit leaders measured?
6. What is the risk tolerance of the company? Understanding the level of risk is very important in making strategy judgments.

7. How does geographical expansion play into the plans?

8. What is the customer's perspective? This is likely to be sourced through marketing.

3.0 Clearly Define Strategic and Tactical Elements

It's important to separate strategic and tactical elements. The boundary between the two is often blurred, leading to wasted time and effort on the part of the storage team (see note 1). From a high level, organizational strategies and business goals form the underlying goals of a storage strategy.

- Strategic storage goals include cost reduction, implementation of environmentally friendly solutions, modernization and the protection of organizational data.
- Tactical elements are the technologies and processes that achieve these strategic goals. For example, these could include the use of Serial Advanced Technology Attachment (SATA) disks in an automated tiered storage array, and the archiving and maintenance of data on tape in two separate physical locations.
- Success indicators need to be agreed on, set and monitored to determine if a storage strategy is being achieved. If cost

optimization is a key strategic component, then cost targets should be set.

Consider the following to discern whether an element is tactical or strategic:

- **Timeline.** Does this element have longer-term business value, independent of technology choice? How would it fit into a hierarchy of actions? Long-term elements are strategic, whereas short-term elements are tactical.
- **Durability.** A task or goal that changes every six months or so is likely not strategic.
- **Business value.** Strategic issues tend to focus less on technology and more on business value, process and people, in that order. Technology is most often an enabler.

Many companies cannot agree on what data should be protected. In this case, the storage team's strategy will prove to be an important guide to help leadership determine business requirements.

4.0 Align Storage Goals With Business Objectives

Align storage goals to business objectives by defining services and service levels and developing service-level agreements (SLAs). SLAs define the deliverables for the storage team and set the "customers" proper expectations with a measure of delivery

Note 1. The Difference Between Strategy and Tactics

Following are two quotations from Sun Tzu's "The Art of War":

- "Strategy without tactics is the slowest route to victory. Tactics without strategy is the noise before defeat."
- "All men can see these tactics whereby I conquer, but what none can see is the strategy out of which victory is evolved."

quality on the service. If SLAs are not aligned with customers' requirements to deliver their business objectives and facilitate day-to-day activities, then discontent will soon set in. The objectives are ultimately linked to price and performance, and this is where the storage strategy and tactics balance the objectives (SLAs, operating-level agreements [OLAs]) with the cost of the tactical solutions implemented.

Services and SLAs can and should evolve over time as business needs change. SLAs also provide a way to measure the progress of the storage team and the quality of the service delivered. Customers should pay for the class and quality of service delivered.

Action Items:

- Formally define services and SLAs, taking input from customers as to their business definitions and SLA definitions, and providing regular review for possible changes and improvements.
- Communicate SLAs and goal alignment regularly to all stakeholders.

5.0 Create a Vendor Selection and Retention Process

The storage team needs to take a comprehensive view of the vendor selection and retention process in a storage technology area. There are various options, each with advantages and drawbacks:

- **Single vendor.** This simplifies staffing issues, because there are fewer sets of tools, application programming interfaces and procedures; little or no interoperability and finger-pointing issues; and fewer mistakes. The downside is that obtaining competitive bids requires a substantial vendor management program.
- **More than one vendor (usually two).** This allows bidding wars for purchases. The disadvantage is that work increases

to deal with the issues outlined above. In addition, if the user is not careful, then negotiating leverage can be lost.

Sometimes, storage decision makers develop an unhealthy reliance on a vendor, believing that vendor to be strategic. This is not necessarily true, and can create wasteful dependencies on a given vendor or service provider. If a technology, product, service or supplier can be replaced within a couple of months by a different, equivalent solution, then it is not strategic. A storage team can test this by asking a vendor to replace another vendor's solution and inquire about the time scale involved. Obviously, the cost to switch must be taken into account. If the costs are too high, then the organization may be too dependent on a solution, product or supplier. If an application or service is tied to or locked in with a specific vendor's technology and exhibits these "sticky" qualities, then it should be reviewed. Although such dependencies may solve specific problems, they often reduce future options and make a storage infrastructure inflexible, as many strategic directions or tactical solutions are tied to a proprietary or point solution.

Action Items:

- Decide how many vendors have products that are appropriate, create a vendor management process, and drive vendor selection based on requirements and desired technology/services.
- Identify strategic suppliers and beware of dependencies. Review "sticky" situations.

6.0 Evaluate Requirements for Management Products and Technologies

Often, what goes wrong in storage is that inappropriate products were selected, or personnel haven't been trained on the tools. The storage team should evaluate and document the requirements for tools and map them to available products so that requirements for

supplementary tools or customization are well-understood. The strategy should also comprehend organizational relationships and the need for training.

Most storage management tools deal with the storage domain, but some vendors try to integrate and overlap server and storage domains into one toolset. Others focus on domains such as VMware deployments or Windows deployments. These approaches can sometimes simplify activities that bridge two different parts of the organization.

Storage demands are application-driven and they can be very different, depending on the type of application. A strategic plan should divide the applications into at least two segments: those that are growing at exponential rates, such as Web applications and e-mail, and those with much more sedentary growth, such as accounting databases, HR and other similar applications.

Action Items:

- Establish centers of expertise to monitor and evaluate new technologies against overall storage requirements and the balance of risk versus payback.
- Carefully assess and plan requirements for management and reporting tools before selecting/purchasing products.
- Monitor new technologies as aids to meeting goals, but create road maps to implement when risks are at an appropriate level.
- Consider fast-growth versus slow-growth applications.

7.0 Use Technologies and Processes to Control Costs

Networked storage, both storage area network and network-attached storage, create opportunities for the deployment of technologies and processes that can dramatically reduce the cost of storage growth. Technologies such as thin provisioning, deduplication and compression are increasingly available for some data types, and tiers offer lower-cost classes of storage by trading off performance. An ideal tiered process should be automated, so as not to require large numbers of staff.

However, while utilization improvement and the creation of storage tiers can provide a one-time reduction, archiving has lasting benefits with high payback. Unfortunately, this is an application-by-application task, involving software and the creation of an archive. How storage is managed, configured and maintained often can cost more than the cost to purchase a product; therefore, it is very important to investigate how storage management processes can be improved to optimize productivity.

Action Items:

- Create and implement a strategy for tiered storage and archiving, but only for your highest-growth applications and where tools exist.
- Implement data reduction techniques, such as deduplication, compression and single-instance store, to contain the cost of disk-based backup and archiving solutions.
- Investigate and implement procedures and processes that increase storage management productivity.

8.0 Create a Data Management Strategy and Set Policies for Smart Archiving

A key building block of a strategic plan is an understanding of the organization's data management strategy. This is a fundamental question, as it will determine how much data your company can manage. Many organizations operate a capture-all policy, but this quickly becomes unmanageable from a cost of data storage perspective and, in some cases, in application response time. Every business has data that should be archived for preservation and future access. Sometimes, legal compliance regulations demand this requirement. Other times, it is driven by the need for good corporate governance. Still other times, it is fundamental to the business's mission (for example, medical records at a hospital).

Archiving brings with it the burden of establishing and maintaining an archive. Although archiving software will likely run in or tap the application server(s) with an agent, few application or system administrators will be interested in responsibilities and processes that last for years.

The storage team should own and manage the physical assets of the archive, with policies set by management and driven by business requirements. Because the data is no longer inside the applications, it may need an owner, traditionally called the archivist.

Action Items:

- Document requirements, strategy and management policy for all data to be archived.
- Try to use SLAs, OLAs, and the recovery time objectives (RTOs) and recovery point objectives (RPOs) definitions to determine archival requirements; incorporate these into a service catalog.

- Work with management to develop policies for the archive, and appoint an archivist.
- The value of data and a thorough understanding of application importance is critical to be able to determine an underlying storage and data archiving strategy.

9.0 Take Inventory of Storage Infrastructure

A storage strategy should take into account a profile of the current installed base. This is especially useful in the creation phase, as it produces a reference point from which IT managers can separate the strategic from the tactical. IT managers may find strategies, goals or tactics that are no longer relevant or required. The inventory should include all aspects of hardware, software, tools, suppliers and people. A full product inventory should also include service costs (products, services and financing) and storage utilization. It's also important to assess your storage infrastructure maturity to plot in progress in an overall strategy plan. A relatively mature storage infrastructure provides services to clients based on agreed-on SLAs.

Action Items:

- Conduct inventory regularly so that it doesn't become a daunting task.
- Consider talking to suppliers to get information.
- Review all applications to determine if storage SLAs, OLAs, RTOs and RPOs have been agreed on and documented.
- Assess the current state of the storage infrastructure as a starting point for a road map to the desired state, how bad or good it is, and what can realistically be done to make it better.

10.0 Project a Timeline to Meet Business and Technology Changes

Typically, a strategic plan projects a time span of three to five years, although this depends on the volatility of the segment in which the organization operates. However, some business requirements are constant, and storage teams should document them. Most organizations' major objective is to remain in business, so storage strategies should outline how they help to achieve this in terms of availability and quality of service. For example, if organizations have a business strategy to acquire or merge with other organizations, then the strategy should specify the ability to scale and be flexible to accommodate large changes. The ability to virtualize storage may be an important component within the strategy. For companies that plan to grow organically, flexibility and the ability to support and incorporate many different types of storage may not be that important, and storage virtualization, which provides the ability to migrate and manage heterogeneous platforms, may not be a strategic direction.

Predicting business conditions, as well as technology changes, beyond this time frame is a statement of a direction, not of strategy, but a good understanding of future conditions may shape and inform the strategy. From a tactical perspective, Gartner warns against embarking on a project that requires more than 18 months to yield a return on investment (ROI). Resources, both financial and human, will determine what can be achieved. Most companies have a mid- and long-term strategic plan, with a focus on a three-year outlook.

Action Items:

- Document business requirements and outline how storage enhances and supports them.
- Embark on projects that have a fairly early ROI.

11.0 Plan for Storage Beyond the Data Center

While a strategic plan will focus on the main storage hubs, it is important to include storage outside the data center. Outlying offices grow and have changing needs. The number of mobile user devices storing organizational or business data is increasing, so the strategy will need to consider this and plan for distributed data.

Action Items:

- Determine how best to manage storage from a cost, performance, service and security perspective for remote locations.
- Determine how storage meets the business continuity requirements. What will be the replication strategies?
- Consider strategies for protecting data for remote offices and mobile employees, who use mobile devices that store corporate data.

12.0 Match Processes to IT

Process plays a vital role in storage strategy and determines how the strategy will be executed relative to other IT plans and, more importantly, to the business itself. The failure to match business and IT processes is one of the most common reasons IT fails to deliver promised services and projects. An IT focused framework such as Information Technology Infrastructure Library (ITIL) is an important element, but it needs to be matched with business processes understanding. To optimize productivity, some organizations determine the storage processes before any hardware or technology selection.

Action Items:

- Review how the project-planning process takes into account the storage strategy.

- Determine whether your organization will make it a mandatory requirement that all departments adhere to the storage strategy.
- Decide and determine the optimum processes to manage and maintain the storage infrastructure.

13.0 Frequently Update the Strategic Plan

Successful storage strategies need ongoing review and an update process that is tied into the planning and review cycle of the business. Technology and practice opportunities will be found more easily if resources on the storage team are charged with the responsibility of tracking the evolution of the industry. However, developing a storage strategy takes time and resources. It may involve legal and administrative expenses, and it could take weeks or even months to collect the input needed to create a storage strategy that reflects organizational goals and is understood by the leadership, IT and business units of the organization (see Note 2).

Gartner recommends that the storage team members responsible for the strategy meet whenever the overarching organization changes its strategy or business model, updating the storage piece accordingly. The group should hold quarterly meetings to review and verify the tactical solutions implemented and propose ways to meet strategic goals.

Finally, proposed changes should comprehend a timeline and road map so that budgets and anything impacted by the changing strategy can react to and take advantage of the update.

Note 2. Constraints That May Affect Your Strategy

Having distilled the input from many areas, it is important to recognize and consider the constraints that may affect your strategy, including:

- Budget
- People/skills
- Risk tolerance
- Legal issues
- Environmental issues

14.0 Create Milestones to Track Progress

Milestones consist of discrete, measurable points that indicate the effectiveness of the strategy. It's important to separate technology milestones from strategic ones, which drive the business value. For example, the implementation of a SAN-based storage system is tactical, but the agility or the ability to quickly scale or reduce storage capacity is a strategic goal that drives and determines which tactics and technologies to implement. Some components of the storage strategy may be simpler and easier to implement than others. For example, data archival requirements can take a long time to agree on, as they are dependent on many external factors, such as legal and compliance requirements.

- *Action Items:*
- Work with the IT and business unit management to develop strategic milestones, and solicit their input in the review process.
- Meet with the IT group to produce realistic strategic, tactical and technical milestones.

15.0 A Checklist of Items to Consider

A large number of factors need to be considered to create a storage strategy. Although the list in Table 1 does not contain all of them, it provides many of the key items, questions and subject areas which need to be considered.

This research is part of a set of related research pieces. See "Storage Strategy: Why You Need One, and What it Should Include" for an overview.

Gartner RAS Core Research Note G00171731, Robert E. Passmore, Valdis Filks, 28 October 2009

Table 1. A Storage Strategy Checklist

Application awareness	Can the TCO of the storage infrastructure be measured?	Determining the key success metric priority list (e.g., cost, agility, compliance, security etc.)
Archiving strategy	Clear separation of strategic versus tactical	Determining what is strategic versus tactical
Are business process framework integrated with IT?	Company awareness	Do processes exist that allow storage strategy to be turned into tactics?
Are the direct costs of running the infrastructure understood (purchase price, service agreements, etc.)?	Comparing before and after	Does a strategic plan exist?
Backup and recovery	Cross-business input	Does an inventory of products and software exist?
Budget sensitivity	Data protection techniques	Disaster recovery
Business process framework integration	Description of strategy	Did we implement an archival plan
Did costs decrease?	How many tiers do you have, and have you consolidated/rationalized?	Rationalization/consolidation
When, why and did we change/reduce vendors?	Identification of strategic suppliers	Risk tolerance
Did storage growth decrease?	Information life cycle management or store everything?	SLA provision goals (includes recovery)
Did reliability increase?	In house or outsource (including cloud)	Sourcing
Did application performance increase?	Is there an information life cycle management or data retention policy?	Staffing policy
Did we have to change the strategy?	IT cross-functional input	Start with the business problem
Was the strategy durable and were the tactics changed to meet the strategy?	Key drivers and inhibitors (e.g., cost, agility, compliance, security)	Technologies: Thin provisioning, deduplication, tiering, compression
Enterprise versus midrange solutions	Management support	Timelines
Frequency of updates	Measuring direct costs versus TCO	To what extent are storage plans integrated with other IT strategic plans?
Geographical considerations	Measuring implementation time	To what extent is the storage strategy integrated with business unit needs and directions?
Hardware, software and process inclusion	Measuring internal and external customer satisfaction	Vendor decisions (including single versus multivendor)
Have SLAs, OLAs, RPOs and RTOs been agreed on for storage?	Measuring part versus all of storage infrastructure	What criteria drive the decision?
How do you manage the supplier?	Storage and data migration	What is its level of sophistication of the strategic plan?
How do you respond to future business requirements?	Milestones	Which application data sets should be considered
How long do you need to keep the data?	Project prioritization	
How many copies and formats?	Protocol decisions	

Source: Gartner (October 2009)

About

Quantum 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, integrated range of disk, tape, 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.

Quantum®

Quantum Corporation
1650 Technology Drive
San Jose, CA 95110
www.quantum.com

