Realtime
publishers

"Leading the Conversation"

Tips and Tricks Guide™ *To*

Creating Business Continuity through Enterprise Storage Solutions

Chad Marshall

Note to Reader: This book presents tips and tricks for seven topics related to business continuity created through enterprise storage solutions. For ease of use, the questions and their solutions are divided into topics, and each question is numbered based on the topic, including

- Topic 1: Securing Availability and Business Continuity
- Topic 2: Maximizing Storage Resources and Solutions
- Topic 3: Information Management
- Topic 4: Cost Management
- Topic 5: Compliance
- Topic 6: Security
- Topic 7: Aligning Storage to Serve Business

Copyright Statement

Topic 2: Maximizing Storage Resources and Solutions

Q2.1: What steps can be taken to ensure a successful analysis of existing storage infrastructure and plan for growth?

A: Ensuring a successful analysis of existing storage and planning for growth is an exercise in capacity and performance management. C&PM is a combined discipline that consists of management processes whose goal is to ensure that IT and related infrastructure continues to meet current and future business needs in the most efficient manner possible.

A successful C&PM program should:

- Help ensure consistent, predictable storage performance
- Reduce failed customer interactions due to performance degradation
- Optimize resource utilization
- Proactively manages risk associated with technology upgrades and growth

Within C&PM, the analysis of existing storage infrastructure and planning for future growth are covered under three areas: C&PM planning, C&PM review, and C&PM forecasting.

C&PM Planning

C&PM planning is the overarching planning process to understand the needs of line of business partners as they relate to the demand for capacity and performance. C&PM planning must be based upon some form of pre-agreed to understanding between line of business partners and IT (or storage managers) on resources and is usually covered under a Service Level Agreement (SLA).

An SLA is a document that clearly outlines the needs and expectations of the business for a particular IT service offering, such as storage. During the planning process, the service delivery manager (or IT manager responsible for the establishment of the plan for C&PM) will work directly with line of business or business technology managers to determine the appropriate and acceptable levels of service that will encompass both the planning and provisioning of infrastructure capacity and availability as well as documenting these needs. For example, a line of business may require that certain data be kept available for 90 days so that anyone who needs access has it immediately; however, following the 90-day period, the data can be stored offline and must be retained for a period of 3 years. Based upon measurements, 90 days worth of online data for this department may be identified as 1TB. This definition is critical in not only understanding the capacity needs of the line of business but also putting a much finer point on the underlying nature of the data, which may be used later in other storage disciplines such as hierarchical storage management (HSM).

C&PM Review

Although the focused reaction of many storage managers may be to cut directly to the "bottom line" and view storage infrastructure needs as a standalone deliverable, determining the current C&PM baseline involves aligning not just storage but also all infrastructure, systems, and resources to an application portfolio. An application portfolio is a way to track an entire application system through an enterprise environment and is one of the most common alignments of infrastructure assets to business needs. Other alignments used for C&PM purposes may include geographical alignments, which are useful when measuring the effectiveness of a site or data center; and platform alignments, which can be used to provide a comparative analysis between different storage or infrastructure platforms within the enterprise.

If you assume the application portfolio approach, your organization may have, for example, a customer relationship management (CRM) system that consists of a CRM database application, an email and newsletter campaign system, a follow-up call-queue system, and various other parts. Together, this system may have been brought into the environment as a total CRM solution; bought and paid for by a single line of business and managed by a single application manager. Although there may be a half dozen or more "applications" in the traditional sense of the word and multiple host devices, servers, and storage resources needed for these applications, the entire portfolio can be managed centrally as a single group with common underlying business drivers. Aligning these individual applications to a portfolio will aide in forecasting by addressing a particular need in terms of multiple applications.

Once application portfolios are aligned to resources, the next step is to gather all the resource metrics for those application portfolios. Resource metrics may consist of mainframe processing utilization, midrange server utilization, storage utilization, network utilization, and so on. Each area of resource utilization has the potential to impact the other, which makes approaching capacity and performance management from a holistic vantage point a necessity. The process of gathering capacity and performance data can be defined as a six-step cycle as illustrated in Figure 2.1.

Figure 2.1: C&PM planning methodology.

The following list highlights each step:

- Data collection—During the data collection process, information is gathered regarding the current capacity and performance. For example, on a storage device, you might gather that out of a 100GB drive, 87.5GB is free (capacity).

- Analysis—The analysis process is arguably the most critical step in C&PM planning because it is at this stage that the data that has been gathered first begins to take a useable form. During the analysis process, data is aligned to application portfolios, geographic regions, server farms, or similar C&PM focus directives to determine a foundational view that is established in the next process.

- Baseline—Based on the analysis, the baseline process will establish a standard base level of expectation.

- Predict—Using input from the analysis as well as information from business volume drivers, the prediction process will derive a prediction based upon known past capacity and performance and current baseline measurements and will predict a future capacity and performance state.

- Validate—Providing backup to the prediction process is the validation process, which measures the difference between past predictions and present "real-world" results to adjust prediction methodology as necessary.

- Report—The goal of the entire exercise is accurately report to all interested parties.

C&PM Forecasting

C&PM forecasting is often the result of a great deal of effort by many parties including Service Delivery Managers (SDMs), business application managers, and storage managers. Through close collaboration between these parties, an understanding of business volume drivers is derived through a process known as business volume driver capacity forecasting.

Business volume driver capacity forecasting is an interactive evaluation that is performed by SDMs and business application managers responsible for a particular application portfolio and is usually performed quarterly to ensure that up-to-date and accurate business volume drivers are being accounted for. When considering storage as part of business volume driver capacity forecasting, the following items are gathered and examined:

- Current mainframe storage requirements
- Current midrange storage requirements
- Current Network Attached Storage (NAS) requirements
- Current Storage Area Network (SAN) requirements
- Current tape/optical/misc. storage requirements
- Forecast mainframe storage requirements
- Forecast midrange storage requirements
- Forecast NAS requirements
- Forecast SAN requirements
- Forecast tape/optical/misc. storage requirements

This list, which is just a starting point of consideration, is usually arranged quarterly by application and in some cases by hostname of the system requiring the resources. Through the collaborative discussions of SDMs and business application managers, assessments are made based upon current business forecasts and their correlation to storage growth—these assessments are designed to answer one critical question: How well does the organization understand how business drivers impact the infrastructure?

If your organization provides health services, for example, it is important to capture what kind of services you offer, to whom, and how those services impact infrastructure requirements. The easiest way to illustrate this is in terms of storage. For example, suppose that your organization is responsible for health care, and you have a team within your organization that does medical imaging. Medical imaging can produce large files that in some cases need to be transported off-site rapidly for evaluation by specialists at other hospitals. For the sake of example, suppose that these images can be broken into three categories of service:

- Image Type 1 requires 10MB of disk space and must be stored for more than 10 years
- Image Type 2 requires 25MB of disk space and must be stored for more than 10 years
- Image Type 3 requires 50MB of disk space and must be stored for more than 10 Years

Once you understand the storage per client, you can make the translation between business and storage requirements. If one of the health care providers from this department were to speak with a service delivery manager and identify a potential to add 500 new patients per year over the next 3 years, you can more accurately forecast and respond to the need. For example, if the "average" patient has five Type-1 images, three Type-2 images, and one Type-3 image during their tenure of visits, and images are only used for 30 percent of all patients, you can apply a bit of math to extrapolate how this information actually translates to storage growth. Not only does this provide a more accurate forecast but also links the forecasting directly to business drivers that can be discussed in business terms and quantified back for the line of business decision makers.

Q2.2: How can I maximize the utilization of my existing storage infrastructure?

A: Maximizing storage utilization is top of mind for many organizations these days. As demand for data storage increases to meet the requirements for compliance, growth, and standardizations and as IT budgets become spread thin as they try to meet these requirements on top of existing storage needs, making the most efficient use of storage infrastructure is central to the overall success of the IT organization. Although "maximizing utilization" sounds really great, the problem many organizations run into is that they fail to both qualify and quantify the statement and therein fall into the biggest pitfall of IT optimization—ambiguity. For any improvement effort to be of any use, goals must be set that are specific, measurable, attainable, realistic, and timely—aka SMART.

Specific

To maintain focus, goals need to be as straightforward as possible with great emphasis given on what *exactly* the end outcome must be. Being specific means to define the what, why, and how:

- What needs to be accomplished? (Reduction of storage growth by 10 percent.)

- Why does this need to be accomplished? (Current forecast analysis shows a 20 percent need by next year but our budget can only maintain a 15 percent growth and areas have been identified as underutilized.)

- How will this be accomplished? (Site a specific action plan that defines who will be involved and what tools will be used.)

Take steps to ensure that each goal you define is not only specific but also as simple and direct as possible.

Measurable

There is an old saying in business that goes "If you can't define what you're doing as a process, you don't know what you're doing." To the same end, at this point of the SMART process, if you can't measure what you're doing, you can't manage what you're doing. The more well defined (specific) your goal is, the easier it is to wrap metrics around the goal. Be certain to pick a goal that can be clearly measured and define how it will be measured, by whom, and at what points so as to be sure that your measurements are accurate, someone is accountable, and delivery is timely.

Attainable

Ensuring that a goal is attainable will include fostering the establishment of finances, nurturing positive attitudes, aligning the necessary skills and abilities, and promoting the goal among peers, supervisors, and subordinates. If the goal being set is too lofty, buy-in from necessary parties will be difficult; if it's too short-sighted, you will have difficulty defining the goal as relevant to the organization. An effective goal will challenge an organization just enough to make people want to be part of the success while avoiding fear of being part of the failure. The drive and thrill of success can be a big motivator.

📖 A common tactic to make large goals more manageable is the implementation of a multi-generational plan. MGPs break down a single linear plan into multiple "generations," with each generation building on the previous to yield greater success over time. If you find your goal seems too large on paper, try breaking it down into generations to see how it might look when brought in stages of manageable "chunks" of time, or, if finance is the concern, over multiple fiscal years.

Realistic

A measure of a goal's realistic potential can be quite difficult to define in tangible terms. Whether something is realistic is more of a gut feeling that is based upon the goal's attainability coupled with your own understanding of the current culture. On paper, a goal may look attainable, but when you begin to consider the people behind the processes, you may discover that some uniquely talented individuals might prefer to remain uniquely talented and not stretch very far to make the attainable possible. To combat this situation and other items that make attainable goals seem unrealistic, it's important to take the following actions:

- Define a clear plan with responsibilities, action items, participants, and goals.

- Relate the plan back to a benefit to each party. Use language that promotes their vested interested such as "This will help you to…" or "By reducing this item you can save…" Drive home those key points.

- Solicit the support of upper management and if at all possible align your project to the associate's individual performance plan.

- Remember to set goals that do stretch but be careful to not stretch to the point of breaking relationships or support.

Timely

It has been said that timing is everything. Consider the last 10 business emails you have sent. Of those, how many would have reflected better or worse upon you if sent a week ago? What about a week from now? When dealing with timing, what you are essentially attempting to accomplish is a controlled rate of perception. If you push an idea or goal too far, too fast you will create too much friction for the goal to flourish. Likewise, if you wait on a goal for too long, it may become outdated, or worse, someone else may implement it!

Once you have avoided the ambiguity pitfall, you will need to set goals that are specific, measurable, attainable, realistic, and timely within your organization to maximize storage utilization. The following list highlights tips to get you started:

- Utilization can sometimes become a difficult item to quantify across multiple storage platforms. Establish consistent standards and policies across all storage platforms to make the process of quantifying utilization attainable.

- Work towards standard and policy compliance through the development of sound procurement and project management processes. Standards and policies are going to do your organization little good unless they're implementable.

- Choose hardware and software vendors whose products and services are compatible.

Q2.3: What are the best practices for storage resource management?

A: Storage resource management (SRM) is the management of physical and logical storage resources, storage devices, appliances, virtual devices, disk volumes, and file resources. To obtain efficient, reliable, and cost-effective SRM, there are several best practices that should be adhered to.

Best Practice 1: Define SRM as an Intelligent Set of Processes

SRM, as a whole, is often a vaguely defined exercise at best. Many organizations understand that they need to do a better job of SRM but are unsure where to begin. By breaking down SRM as a series of processes, you can begin to take control of the environment. The steps to get control of SRM can be broken down into four major process areas (see Figure 2.2), which are sometimes referred to as Intelligent Storage Management (ISM).
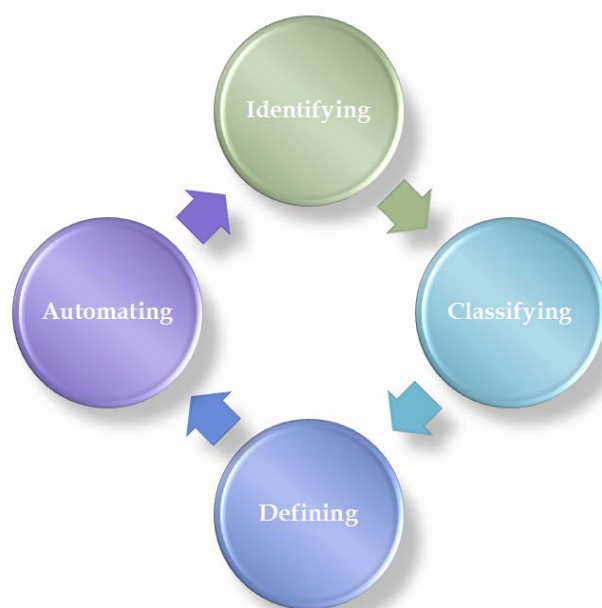

Figure 2.2: Intelligent storage management cycle.

- Identifying—During the identification process, all assets (servers, storage components, network devices) as well as all applications and data that reside or has cause to transverse the storage infrastructure (such as transaction data) is identified.

- Classifying—The classification process categorizes information based upon its importance to business and should include information about who created the data, why it was created, how it relates to business, and how it relates to any specific regulatory compliance concern.

- Defining—During the defining process, the definition and structure of information and its supporting environment is undertaken to establish a series of repeatable processes to ensure information is managed and protected consistently. This definition may include the type of security required or how the data is accessed to ensure information is secure and continuously available.

- Automating—During the automating process, steps are taken to automate as many tasks as possible, which contribute to resiliency by automating reactions to adverse situations and minimize the risk of human error.

🖉 What is the difference between SRM and Information Lifecycle Management (ILM)? ILM is a sustainable storage strategy deployed to balance the costs of storing and managing information in direct relation to the information's business value. SRM is the tactical management of resources to meet ILM goals.

Best Practice 2: Agree Upon a Standard Level of Service

Managing storage resources is as much about the actual storage resource and information life cycle management as it is about meeting the expectations of your storage consumers. Defining that level of expectation and service should be performed by clearly documenting and reviewing requirements and service levels with your line of business partners. Developing a Service Level Agreement (SLA) should be exercised as an SRM best practice.

The IT Infrastructure Library (ITIL) defines processes for obtaining SLAs as part of a larger Service Level Management (SLM) framework. SLM provides for continual identification, monitoring, and review of the levels of IT services specified in the SLAs and ensures that Operating Level Agreements (OLAs) are in place with internal IT support providers and external suppliers in the form of Underpinning Contracts (UpCs). The process involves assessing the impact of change upon service quality and SLAs. The central role of SLM establishes it as a central point for metrics to be identified and monitored against a benchmark.

🖉 What's the difference between an OLA and an UpC? An OLA sets forth a standard, agreed-upon, level of service between two internal parties. For example, an OLA may exist between an internal database management team and the Windows server team that supports the database servers that sets forth a standard, agreed-upon, level of service. In contrast, an UpC is an agreement between an internal party and an external vendor. If the team responsible for servers has a contract with a large vendor to provide on-site support, this agreement would typically be captured as a UpC.

For example, your internal technology organization (Chief Information Officer and below) may establish SLAs with your various line of business partners that clearly set forth expectations of your IT organization. That overarching SLA may mandate that a particular application system be available for production use 99.999 percent of the time during a specified timeframe (Monday through Friday, 9am to 5pm local time). Internally, to meet the needs of this SLA, your team may enter into an internal OLA that addresses the storage availability for the storage resources aligned with that particular application. What has been established through those agreements are metrics that allow your organization to track its ability to meet the needs of your business partners and provides accountability, in the form of Service Level Reporting, which will be discussed in Best Practice 4.

Having this agreed-upon understanding, end-to-end at the line of business customer level will not only help to ensure that appropriate levels of service are met and maintained but also provides storage resource managers with a more succinct view of how their storage infrastructure is being utilized and plan accordingly.

Best Practice 3: Continuously Monitor

Most IT professionals acknowledge that monitoring needs to occur, but they're often uncertain which, of any number of software tools, is best suited for the task. Although software does make monitoring easier, it should not be sought out as a single-point solution that will solve the mystery of monitoring.

Monitoring actually begins with the development of a monitoring approach that defines the scope, methodology, and processes to be followed to continuously monitor your storage resources. During this process you should identify:

- Types of systems that need to be monitored

- High-level methodology to be used to gather the monitoring data

- High-level process to be followed for monitoring

Once the groundwork has been laid, the next step is to begin to define what specific monitoring data needs to be collected and how that data is to be collected. For example, at a high level, you might have identified that you need to monitor storage resources on NAS devices; and now at this level, you begin to articulate what information you need to capture from those devices and how you will gather that data. Once this information is captured, the next step is to begin to identify how to collect the data. The 'how' may come in the form of internally developed or vendor-supplied software.

Although internally developed software has the advantage of being customizable to meet your specific needs at a very granular level, commercial software is much more cost efficient and often has the advantage of being cross-platform compatible, meaning that it can monitor data relevant to SRM on multiple platforms (whereas internally developed or "home grown" software applications often require laborious development to meet the need of each type of storage device within the enterprise). Look for software that is storage vendor agnostic and provides the capability of centralized monitoring, analysis, and automation of storage management operations. A well-designed SRM monitoring solution will also include the capability to automatically react to predefined conditions by invoking predefined actions, and in doing, so will avert system downtime and ultimately contribute to the integrity and availability of the storage infrastructure.

Best Practice 4: Maximize Storage Utilization

Consistent standards and policies should be applied across all storage platforms and locations, and you should continually strive to maximize existing hardware, software, and staff time. Doing so will reduce the need to purchase and support multiple storage products by making the most efficient use of your existing storage infrastructure. Reclaiming storage space can be accomplished by ensuring that stale or obsolete data is promptly removed and by implementing application filters to prevent non-business files, such as music and games, from being downloaded.

Best Practice 5: Reporting

SRM reporting should include provisions for reporting on domains, servers, arrays, volumes, files, and owners. A useful SRM software suite will also have the capability to build and customize your own queries. In addition to internal reporting, you should consider developing service-level reporting based upon your SLAs. Such reports may include data about the number of incidents per month within the storage infrastructure, the incident severity and applications, or lines of businesses, impacted by the incident.

Q2.4: What is the difference between capacity and performance management and how can I leverage them to maximize storage resources?

A: Capacity management and performance management are often viewed as a combined discipline that consists of management processes whose goal is to ensure that IT and related infrastructure continues to meet current and future business needs in the most efficient manner possible. Capacity management focuses specifically on the amount of resources currently available, able to be made available, or forecasted as a need to be available for future use. Performance management conversely focuses on the actual performance of the assets providing the capacity. For example, a 2TB storage array may have a great amount of capacity but if the read/write time is degraded due to fragmentation, or some other manageable cause, the benefit of the array's capacity may be severely degraded or lost to some users entirely.

Capacity & Performance Management (C&PM) as a combined discipline is defined in sub-disciplines, as discussed in Q2.1:

- C&PM initial planning—Often the most challenging discipline to get under control because of its breadth, C&PM initial planning involves the gathering and review of all capacity and performance data associated with IT resources. During C&PM initial planning, the IT or storage manager will gather and then review capacity and performance of IT resources with the goal of ensuring that resources are available and capable of meeting agreed-upon workloads. The output of the planning stage is a plan to improve capacity, performance, or both by focusing C&PM reviews, providing more accurate forecasting, reviewing resource availability, and ensure accurate monitoring and reporting of C&PM data.

- Current C&PM review—C&PM reviews are conducted to analyze current workloads and any past forecasts that lead up to delivery of the present state of service to determine whether the current environment meets the current need. Gathering of any past forecasts enables the IT or storage manager to review the effectiveness of the previous C&PM forecast and make adjustments to the methodology. For example, if a year prior to the C&PM review, a C&PM forecast was done focused on storage for a particular application whose primary storage driver is the number of "customers" and their estimated storage needs, and that forecast suggested that each "customer" needed 10MB of storage for their records, a review at this point would validate whether 10MB was too much storage, too little storage, or accurately represented the need over time.

- C&PM forecasting—Building from a current C&PM review, C&PM forecasting specifically focuses on forecasting future capacity and performance needs. During this stage, information from consumers will be gathered and correlated to capacity and performance as primary drivers, a process that is sometimes referred to as business volume driver capacity forecasting. Each business driver is assessed for its positive or negative impact on storage capacity and/or performance. Building from the previous example, a "customer" was identified as a business volume driver. Other business volume drivers may include "transactions" or, in the case of a mortgage company, "loans". Understanding how each business driver impacts capacity and performance provides a direct link back to the business and ties infrastructure costs to a linear model driven from the same drivers that impact the business. If "customer" is the driver and the business exceeds its growth expectation of "customers" by 200 percent, it's much easier to rationalize the storage or infrastructure expense when both you and your business partners are forecasting growth on the same level and you're speaking the same language.

- Monitoring and reporting—Monitoring and reporting of present capacity and performance is an important pro-active/reactive health measure. Through proactive monitoring, C&PM data is collected for use in current C&PM reviews and to validate forecast estimates and growth patterns. Reactive monitoring and reporting ensures that when infrastructure becomes degraded or reaches a predefined threshold, the proper parties are informed so that they may take corrective action.

- Resource availability—Ensuring that resources are consistently available to meet the demands of the environment is covered is a critical portion of C&PM. Generally, the larger the environment, the more complex resource availability becomes as infrastructure is compounded in larger environments, adding to the overall complexity. Unlike monitoring and reporting, which is focused primarily on the storage or infrastructure component itself, resource availability as a discipline takes into account all the other aspects that come into play. For example, if a network component leading to a SAN within a data center becomes degraded, mechanisms in place to monitor resource availability may determine the potential for a "bottleneck" or degraded performance and inform the IT or storage managers to take corrective action.

Throughout all of these sub-disciplines, it is important to recognize the synergistic relationship between capacity and performance. A great amount of capacity does little if its performance is so low it cannot be effectively utilized, and a great amount of performance does little if the device doesn't have enough capacity to meet current needs of the organization.

> Goals and metrics to get your organization started with C&PM are provided in the Control Objectives for Information and related Technology (COBIT), which is freely available from http://www.isaca.org.

Q2.5: What is "sprawl" and what kinds of technologies and strategies exist today to minimize its impact to the infrastructure?

A: The term "sprawl" has been adopted from its use in land management where it is often employed with a negative connotation to refer to undesired population growth resulting primarily from irresponsible development practices. In the storage management context, "sprawl" is quite similar in that it often refers to irresponsible development, design, adoption, or implementation of storage solutions, the result of which usually leaves the storage manager with a disparate and difficult-to-support environment consisting of incompatible "standards" and architecture.

One of the most important, though often-overlooked, aspects to consider when tackling the problem of storage sprawl is the cause. Many IT and storage managers will too easily classify the problem as one of necessity that occurred over time, when in actuality, sprawl has been allowed entry into the environment not through necessity but rather through poorly documented (or non-existent) storage and procurement strategies. The first step in combating sprawl is to take a step back and correct the situation that allowed it to occur in the first place.

Define a Strategy

Many IT managers consider "strategy" to be the responsibility of the business they support—after all, business should be the driving force behind technology growth. However, developing a storage strategy that complements the business strategy will do a number of things to benefit both the storage infrastructure and the business:

- Increase communication between line of business stakeholders and technology providers

- Demonstrate commitment to line of business goals

- Align technology to business objectives

- Provide greater visibility of business decisions that drive technology needs

Consider your storage solutions strategy as the high-level plan to meet the needs of your line of business partners through your storage infrastructure. It is important when defining strategy to step outside of the storage management arena and partner with internal and external stakeholders and to gain as much industry knowledge as possible. Although it may not always be possible to draw a direct correlation between business strategy and storage strategy, work to remain complementary. In the end, you might discover business drivers that will work as justification to meet internal goals, such as a business need for data analytics and an internal desire to consolidate the storage infrastructure.

Take a Stand—Standardize

Once a strategy has been put in place that complements business needs, it's time to align storage to as much of a standard platform as possible. Standardization can often be received, inaccurately, as a limitation of technology resources, which should not be the case. What standardization should mean is that, when possible, every attempt should be made to use a standard storage offering. Aligning storage offerings to the needs of the business is an exercise in understanding business needs, but once fully understood, storage offerings can be maximized to meet needs and accurately forecast to meet growing (or declining) business requirements.

Standardization should result in the creation of two processes in larger organizations' governing bodies.

- Storage roadmap—A storage roadmap is an agreed upon plan for storage offerings on a time scale (typically 3 years) that clearly shows which technologies will be considered as adopted, approved for use, sunset, or retired.

- Storage review board (SRB)—An SRB defines a storage review process to be used within an organization when assessing storage additions or modifications to existing storage infrastructure. The SRB should have the authority to make exceptions to existing standards to meet business needs and meet regularly to discuss exceptions being raised. The SRB may also be a subset of a larger IT review board, such as an Application Review Board, but should consist primarily of storage key decision makers who are fully versed in storage strategy.

The decision of which product or service to standardize can often be difficult, so try to allow for as much time as necessary to make this decision and weigh as many stakeholder comments as possible. Remember that standardization does not mean that no other storage products or services can be used; it simply means that your organization is attempting to align storage to its needs to the greatest extent possible. Exceptions can and will exist and in those cases in which an exception must be made, other techniques can be used to minimize impact. A support team may, for example, need to offset the cost of additional employees specifically skilled to support the non-standard environment—a cost that should be directed back to the line of business generating the demand.

Stay Standard Neutral

Standardization is a great way to start to bring some order to the storage environment, but there are going to be documented legitimate cases in which storage needs do not align to standard storage offerings. One way to offer some protection against such a scenario is to stay as standard-neutral with storage management solutions as possible. Various storage management software vendors offer solutions that are focused on the storage (rather than the storage vendor) and make storage management in a disparate environment an achievable goal.

Q2.6: What steps can I take to consolidate storage management?

A: Consolidating storage management can be broken down into at least three major categories. First, you must understand the larger consolidation picture, holistically, and the four stages of consolidation. You can then examine the benefits of proactive Capacity and Performance Management (C&PM) and the benefits of introducing hierarchical storage management (HSM).

Step 1: Understand the BIG Consolidation Picture

To consolidate storage management, it is important to first understand the four stages of consolidation:

- Logical consolidation—Logical consolidation is the centralization, or unifying, of management of IT resources and is commonly referred to as unified management. This is the most common and cost-effective area of consolidation and the one that most directly aligns to consolidating storage management.
- Centralized consolidation (co-location)—Centralized consolidation, sometimes called co-location, is the physical consolidation of servers, storage devices, or other infrastructure components to a central location—usually a datacenter.
- Physical consolidation (compiled workload)—Physical consolidation is the compiling of workload from multiple servers or storage devices onto a single server platform. In terms of storage, this may be the equivalent of replacing multiple independent RAID-5 storage arrays on many separate servers with a storage area network (SAN).
- Operational consolidation (application consolidation)—Operational consolidation is the consolidation of multiple platform operating systems (OSs) and application systems onto a single server. In terms of server consolidation, operations consolidation occurs when a server serves more than one purpose. Given the heterogeneous nature of storage, application data is usually not as segregated by storage device and needs to be operationally consolidated.

Within this four-stage model, consolidation of storage management will often be most heavily directed within the areas of logical and physical consolidation. Logical consolidation would include any hardware or software mechanism that allows for the unified management of multiple storage systems, such as software that allows for the central management of all SAN, NAS, and other storage devices; whereas physical consolidation reduces the number of physical storage assets to be managed, thus simplifying the overall management process. Two other focus areas for consolidation are HSM and data replication.

Step 2: Practice Proactive C&PM

C&PM is the ideal tool to identify and correlate opportunities for consolidation in terms of real business benefit because the processes that drive C&PM will:

- Instrument all systems for data collection—This data can then be used to identify opportunities for storage consolidation.

- Map servers and storage devices to application portfolios—Examining storage from an application point of view ties the value back to the line of business paying for the storage.

- Create capacity baselines—Baselines are the foundation for comparative measurement and can be used to compare current capacity with expected capacity. When capacity is shown in excess, a case can be made for consolidation.

- Forecast capacity utilization—Forecast results that indicate capacity surplus can be used as a basis for a consolidation initiative.

- Validate the forecast—Delivers some level of assurance that forecast capacity utilization is accurate within a certain percentage.

Once complete, the C&PM data can be used continuously to identify opportunities for consolidation throughout the storage environment.

Step 3: Implement HSM

High-speed storage devices, such as hard disk drives, are an ideal place for storing just about everything from a strictly performance standpoint; however, the reality is that hard drives are more costly than offline storage media and consume power. Magnetic media, optical disks, and other offline storage media offer a cost-effective solution but disconnect the data from being readily or rapidly available. Determining what data is worthy of inclusion on online storage vs. what data can be sent to offline storage thus becomes the big question. HSM aides storage managers in answering that question.

HSM is an automated and policy-based storage management mechanism that uses "data types" to assign data to specific kinds of storage media. Data types are determined in a number of ways such as data priority, need for accessibility, and need for retention. HSM lowers the total cost of ownership (TCO) for storage by organizing data to be stored on the most efficient storage media possible. HSM also automates the management of data retention schedules and provides for automated data reclamation.

Q2.7: Do any opportunities exist to increase scalability without significantly increasing costs?

A: Although there is no one-size-fits-all solution to guarantee storage scalability for every need in every environment, enterprise, or circumstance, there are a few steps that can be taken to increase scalability without increasing costs.

Step 1: Establish and Maintain Storage Standards

A standard is a document that defines the specific use of technology and the manner in which it is used. The implementation of standards within an organization provides for uniformity throughout the enterprise. For example, your organization may have recognized the use of a specific kind of SAN architecture as a standard to provide for uniformity of SAN storage architecture throughout your organization.

Implementing standards today will help to control the purchases of tomorrow and combat storage "sprawl" (a concept covered in Q2.5) that greatly degrades an enterprises ability to become "scalable."

Step 2: Begin to Combat Storage "Sprawl"

Removing absolute or incompatible storage elements from the environment will pay dividends in scalability by moving the data contained therein to storage devices that can be more readily, and centrally, managed. The act of fighting storage "sprawl" begins by selecting a single target. Start small. Pick a few incompatible or obsolete devices and begin to migrate them over to newer storage media as they become available. When combined with a C&PM exercise, this tactic can make for a powerful business case from both a line of business perspective and an IT/storage management perspective.

Step 3: Establish Unified Storage Management

Unlike server scalability, which falls into a silo and is traditionally managed more as physical assets than as virtual commodities, storage is a bit more virtual. After all, from a user or even an application perspective, there is very little that might restrict the physical location of the storage component. In fact, most people are completely oblivious to the physical location of their mapped drives or where data is physically stored. From this vantage point, storage is a perceived as a homogenous entity, while IT and storage managers, much to the contrary, realize its heterogeneous nature. Utilizing software that will enable to you to manage disparate hardware and standards through a common interface will further the homogenous illusion to the administrator and greatly simplify management—and simplified management enables scalability.

In addition to these three steps, you should begin to gain support for those ideas that you know are going to cost money. HSM begins with a conversation, as does C&PM and all of the other topics discussed within this guide. To gain support for an idea, it must first be discussed, and discussions often lead to actions. Sowing seeds of information and understanding now will pay off in the end.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.