The Complete Reference

Storage Networks

Chapter 24

Capacity Planning

Planning for storage capacity requires the analysis of future end-user demands and the effective management of existing storage resources, both of which support the ongoing processing demands of the business. Provisioning storage can be as simple as adding new storage devices to existing application servers, or as complex as the calculation of new storage demands for supporting a global Internet services application. Regardless of complexity or application, storage demands are end-user driven and reflective of the dynamics within the economics of application usage. For example, an upward trend in sales orders may generate an increase of 100GB of storage to the online application to handle the new business. New customers may almost always spark an increase in storage capacities as the systems expand to handle the customer and related files.

This expansion, and sometimes contraction, of user demands creates the necessity for planning the processing configurations necessary to meet these demands. This chapter brings our cycle of systems management disciplines to the starting point (see the following illustration): the beginnings, development, and/or evaluation of a capacity plan. Capacity planning for storage networks is a subset of overall storage planning, which supports a macro capacity plan for the enterprise, or at least the major components of the data center computer configurations. Consequently, our discussions within this chapter will reflect the support of the entire data-center infrastructure, while articulating the value of developing a capacity plan for storage network infrastructures.

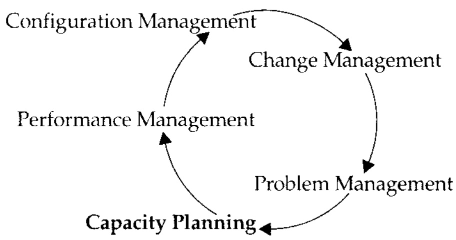

Over time, data centers develop activities and practices necessary to keep up with computing capacities. This is generally referred to as the capacity plan. These practices and activities have matured and congealed into areas of specialization driven by the diversity that exists even in the most conservative IT organizations. These activities are generally divided along physical areas of responsibilities, as well as technological spheres of influence. Figure 24-1 indicates both the physical division of responsibilities within most data centers, as well as the expected technology segregation. This shows applications, systems, and networks as the dominant areas of responsibility. However, supporting the end-user community is generally an operational area that supports end-user computing and help desk functions, and which acts as a conduit of information driven by day-to-day operations of the major areas of specialization (for example, applications, networks, and systems).

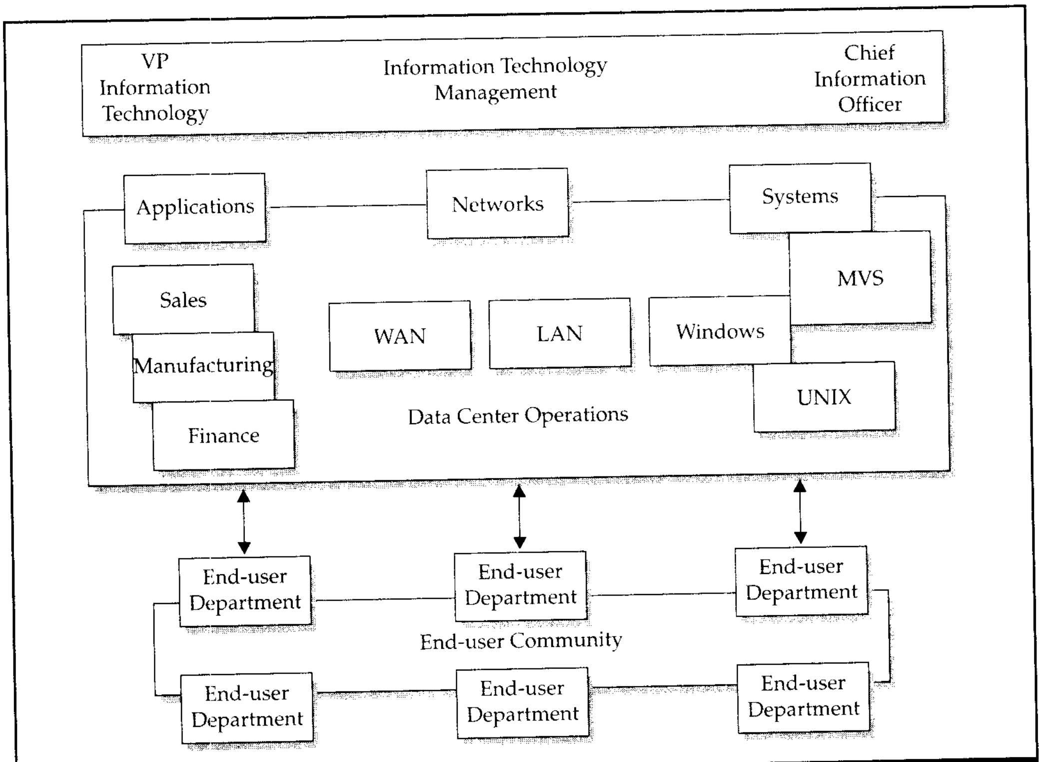

Figure 24-1. *The traditional division of responsibilities within the data center*

Hidden within these organizational divisions are the sometimes invisible spheres of technology influence and expertise. These are characterized by the segregation of the application "elite" and the "back room" systems mechanics. This forms the major factions that deal with capacity planning activities. Given that capacity is driven by end-user demands, the application elite is called upon to use their system's analyst skills to communicate end-user future plans in terms that can be translated by the system's mechanics. Within this milieu of communications, the operational front liners form the ad hoc reality-based view of capacity requirements.

Consequently, one of the major segments of this chapter is the discussion of how user requirements can be interpreted properly and used as valid input into the storage planning process. As with any human communications endeavor, this requires a bit of diplomacy and tact to accomplish. Given that storage, and in particular storage networking, is often viewed as a third-level component within the enterprise capacity plan, care and understanding of the communications process provides a certain amount of value in accomplishing a successful storage networking capacity plan.

The importance of integrating storage into a larger plan cannot be overlooked. The need to articulate and justify complete resources (this includes hardware/software, facilities, and the human element) can make the initial enterprise capacity plan (that includes storage) a success. The reason is that storage has often been viewed as a commodity and is purchased in a manner similar to all processing commodities like memory, end-user software, and printers, by cost. However, the storage network changes that significantly, by offering a complete infrastructure that includes the support of multiple processing server configurations. Storage networks can hardly be seen as a commodity, regardless of what vendors want you to think. On the other hand, NAS, although easy to slip through the official capacity planning exercises, often has extenuating effects on other supporting infrastructures within the data center, such as the network.

Finally, the justification of a SAN or NAS configuration based solely on consolidation of storage, servers, or recovery purposes does not make a capacity plan. These can be characterized as internal user requirements that may or may not justify the use of either storage network solution. This is not to say that these are superfluous or self-indulgent requirements; however, they must relate to the business supported by the data center. Justification on a cost basis, which is usually the case for consolidation and recovery strategies, needs to be balanced by a long-term plan to allow the storage infrastructure to grow along with the storage demands of the business.

Also, the establishment of a plan does not dictate its success. The data center must manage the plan. This is especially important with technologies that are volatile, support dynamic growth, and which are critical to an application's success. This is what storage networks are, and that's why they need a storage networking capacity plan of their own. Once integrated into the enterprise plan, and having the confidence of data center management, the storage plan can play a vital and productive role in the long term.

Developing a Plan for Storage

Storage provisioning begins with a sound capacity plan to establish a reasonable configuration that supports the business applications and constituency of users—both end users and data-center users. Thus begins a chicken-and-egg scenario. It also provides the justification for things like beta installations prior to committing new technology platforms like SAN and NAS configurations to production environments.

To calculate and commit to availability metrics, the following guidelines should be considered in establishing a reasonable plan to support SAN or NAS configurations.

Analyzing End-User, Application, and Internal Systems Requirements

- Develop a method of communications with end users, applications designers and developers, operations, and systems infrastructure colleagues.
- Define a capacity storage metric and level of precision. This could be capacities in number of bytes with a precision to the nearest megabyte, or it could be a more esoteric work unit such as an application work unit calculated as 1 gigabyte per unit, or a relational database unit such as 1 megabyte per unit. Whatever you use, keep consistent with the metric and level of precision. Keep in mind that you should allow for incongruities in media, which may be reflected in disk and tape densities, sector sizes, and blocking architectures.
- When collecting storage capacity requirements, remember to calculate both end-user aggregate data and the necessary overhead to support the data. For example, if end-user data for a database is 100GB, then the total data requirement will be 100GB plus the necessary overhead of database system, index, and temporary tables. Depending on the type of application, this could go as high as 50 percent of the aggregate user data storage requirement, rendering our example to 150GB required for current database physical storage. Other application systems such as e-mail, ERP, and CRM all have additional storage overhead requirements that are important to consider along with the aggregate user data.

Planning and Establishing Adequate Capacity Based Upon I/O Workload Estimates

- An important consideration to storage capacity is the processing requirements that together with capacity requirements drive the configuration elements necessary to meet end-user demand. In order to accurately estimate these requirements, its necessary to translate user requirements into I/O workload estimates (see Chapter 17).
- Using I/O workload identification recommendations (again, see Chapter 17) the application workload types begin to distinguish themselves. Consider size versus access as the applications demonstrate their workload characteristics such as OLTP, batch, messaging, and combinations of these. The size factor, as depicted in our discussion about capacity, must be considered along with the calculation of the I/O workload factors. This allows for the initial elements of the configuration to start to take shape.

- Recovery factors also influence the configuration elements. My I/O workload guidelines provide a recovery factor for calculating the configuration elements (see Chapter 17). However, this calculation should reflect the necessary insurance required for particular data-center business continuity requirements (see Chapter 21).
- Although a subset of recovery, the redundancy factor must not only be taken into account for recovery factors, but also for security and individual failover. Largely driven by system failover requirements, the redundancy in both NAS and SAN can be implemented within themselves, and thus can be a substantial factor in configuring redundant systems. Therefore, they should be considered as separate items even though they're related to recovery and business continuity at a macro level.

Establishing an External Capacity Driver Matrix

- Application requirements drive most of the capacity requirements; however, they provide little if any insight into the storage configuration to support the business application. By developing an external capacity driver, you have identified the major influences that drive both capacity and performance.
- System requirements drive the rest of the capacity requirements with overhead to OS, subsystems, and network requirements. However, these spheres of technology also exert influence on the configuration they feel are most appropriate to the processing of the business applications. By developing the external capacity driver, you have also identified the influences on the storage infrastructure.

Establishing a Reporting Mechanism

- Once the end user, application, and systems requirements are translated and calculated into requirements for a storage configuration, develop a reporting mechanism on the activity of their storage usage. The system must be able to be as proactive as possible in order to show trends that are outside the requirements, and to compare them against the storage capacity plan.

Managing to Storage Capacity Plans and Established Service Levels

- The establishment of the reporting mechanism begins the management of the plan. Given the translation of requirements to real storage resources, the users must be held accountable for their estimates. However, given the volatility of the technologies and the changes in business climates, things will change. Therefore, it's important to also establish an effective communication system track requirements to actual usage without retribution (see establishing a non-aggression pact).

- The User Non-Aggression Pact: It is always helpful to reach an agreement with a user community that accounts for the unexpected. This allows both parties to agree that the unexpected does happen, and should it occur, it should be resolved with both mutual understanding and changes to the existing plan. This can be critical given the volatility of SAN installations and NAS data volume requirements.
  - **Unexpected Requirements** The most common set of unexpected circumstances is the unforeseen application that requires support but which has not been planned for. In terms of affecting the SAN installation, this can be costly as well as disruptive, given the scope of the new application storage requirements. Within the NAS environment, this is one of the strengths NAS brings to the data center. If the requirements can be handled through file access, NAS performance, and capacities, then the NAS solution can be used effectively during these circumstances.
  - **Unforeseen Technology Enhancements** This is the most common circumstance when initial SAN design proves insufficient to handle the workload. The outcome of retrofitting the SAN configuration with enhanced components is additional cost and disruption.
  - **Mid-term Corrections** It's likely that any enterprise storage installation will experience either one or both of the preceding conditions. Consequently, it's extremely important to build into the user agreements the ability to provide mid-term corrections that are an evaluation of the current services and corrections to requirements in order to continue to meet the committed services.

The Storage Analysis

Analysis of storage is divided into two distinct parts: new storage demands and existing storage allocations. Although these two different activities culminate in the same place, which are reasonable configurations that support end-user demands, the reason they are distinct is the following. The demands for new storage provide an opportunity to consider alternative methods to meeting the demands, rather than extending the current configuration. This is most appropriate in storage networking given the sometimes overlapping solutions that exist between NAS and SAN solutions. It is also appropriate when considering moving from direct-attached configurations where the storage demands are not closely tied to existing configurations.

Leveraging a Storage Upgrade

When considering a storage upgrade for capacity and performance, the analysis of a storage network should be part of the storage capacity planning activities. However, this does provide a challenge to the existing server infrastructure and installed storage components. Moving away from any of these requires additional hardware, software, and training. It also requires concurrence from the external drivers, such as applications and systems. Likely, each of these may be resistant to change and the ability to rely on information from the external driver matrix will assist in the justification of the "what's in it for me" scenario.

Analysis of existing storage can provide greater justification into storage networking solutions given the scalability limitations within the client/server direct-attached model (see Chapters 1 and 2). More often than not, the move to storage networking provides a longer term solution in supporting increased user demands. Using the previously described user requirements translations and I/O workload calculations, the justification can prepare the way for the necessary increase in expenditures as well as showing the short-term fix that adding storage to existing server configurations will have.

Establishing a New Storage Network

Analyzing new storage demands provides an opportunity to leverage a storage network— driven by user requirements and I/O workload analysis, the justification can be compelling. New storage capacity can be depicted in terms of scalability of capacity and performance, but also in its ability to consolidate some of the legacy storage into the storage network. This provides the first articulation and integration of internal consolidation factors that are so popular in justifying SANs. However, the same can be said for NAS devices if the storage and application characteristics are justified in this solution.

Working in conjunction with your systems colleagues, there can be real synergy in establishing a storage network strategy. First, is the consolidation of servers. That, in itself, is a large cost factor reduction in overall systems responsibility and administration. This will be augmented by added savings of OS and application license fees associated with multiple servers with direct-attached strategies. Finally, there is the added benefit of managing fewer server entities and the processing consolidation that occurs with the collapse of application processes into a single, albeit larger, server. These cost savings start to mediate the increased costs associated with a storage network solution.

A hidden benefit to systems administrators is the performance and problem management factors that come with storage networking. The consolidation of servers and the collapsing of the storage arrays translates to less network connectivity, minimum servers to manage, and consequently fewer things to go wrong. Establishing an agreed upon and reasonable metric for this allows a quantifiable benefit to be monitored when adopted. In other words, if putting in four NAS devices can collapse file servers on a 20:1 basis, then the quantifiable benefit will be losing 20 general-purpose servers for every single NAS device. If that were the case, then once the plan is implemented, the redeployment or retirement of the 80 servers, to use our example, would create excellent credibility for the plan and storage network.

Tools for Storage Analysis

Storage networking can be analyzed in two ways: the physical capacity and the performance of the configuration. Just as direct-attached storage configurations are monitored for storage allocation, usage, and access, storage networks need to provide the same information. The problem comes with both the lack of tools that account for multiple usages of networked storage arrays and the immaturity of system tracking databases that provide historical data. Performance monitoring in storage networks provides the same challenge regarding tool deficiencies and lack of historical data collectors.

There are several storage software choices when considering capacity and access monitoring tools. It is beyond the scope of this book to analyze or recommend any of these tools and should be a data-center specific choice dependent on the specific needs of the entire storage infrastructure. These tools fall into two categories related to a wide variety of tools known as storage resource management tools. Subcategories include the quota management and volume management tools.

Storage Resource Management

Quota management tools provide a mechanism to assign storage capacity quotas to end users or specific applications, or a combination of both. These provide a safety net in terms of storage usage and the potential of errant users or applications to utilize the majority of storage capacity in an unregulated environment. The quotas are generally set by administrators, either from a storage or systems perspective, and managed from a central location.

The majority of tools in these categories provide some level of reporting on storage utilization. The more sophisticated tools provide views from several perspectives: by end user, by application, by logical volume, or by logical device. These tools work well with file-oriented storage implementations; however, they become problematic when attempting to understand the storage usage of relational databases or applications employing an embedded database model.

It's important to note that in these cases, and given the installation of databases within the data centers (which is likely to be large), the monitoring of storage utilization and access needs to rely on the database monitoring and management tools. This provides another important element within the external capacity driver matrix (see establishing an external capacity matrix), which is the influence of database administrators, designers, and programmers. Given the transparent nature of the relational and embedded databases to the physical storage, the usage of these application subsystems needs to be managed in conjunction with the database expertise. However, it's also important to become familiar with these tools to understand the underlying physical activity within the storage infrastructure.

Volume Managers

Another category of tools that can provide storage analysis information are volume managers. These tools provide the ability to further manage storage by allocating the physical storage in virtual pools of capacity. This allows file systems and applications to access particular volumes that are predefined with specific storage capacity. This can be extremely valuable when allocating storage for particular applications that have specific requirements for storage needs and may become volatile if storage becomes constrained.

Like quota management tools, most volume managers have reporting mechanisms to track both usage and access. They also provide an important function and level of detail that enhances both performance analysis and problem determination. Volume managers work in conjunction with storage controller and adapter hardware to separate the logical unit (LUN) numbering schemes used by the hardware functions. As such, they provide a physical-to-logical translation that becomes critical in understanding the actual operation of a storage configuration. Given the complexities of SANs, these functions can be extremely important in problem determination and monitoring performance.

We have used volume manager examples in many of the figures within this book. As Figure 24-2 shows, they add relevance to naming storage devices within a configuration. As seen in this figure, a typical storage area network, prod01, prod02, and prod03 are disk volumes that contain the production databases for the configuration. The storage administrator through the services of the volume manager software assigns the specific volume a name. The volume manager, meanwhile, manages the storage pools, prod01, prod02, and prod03, transparently. Working in conjunction with storage controllers and adapters, the volume manager works to translate the LUN assignments within each of their storage arrays. In viewing the actual physical operation of the storage configurations, one must understand and look at the LUN assignments and activity within each storage configuration.

The quota and volume management tools provide a snapshot of storage usage and access. However, looking for a longer term historical usage becomes problematic. One of the continuing difficulties within the open systems operating environments is the lack of historical collection capabilities, although this situation continues for a number of reasons, such as disparate processing architectures, box mentality, and distributed computing challenges. Each of these provides a roadblock when systems personnel attempt to gather historical processing information.

Collecting Storage Information

Disparate operating systems provide the first problem, as the major contributor to the differences between vendors, but especially between UNIX and Windows operating environments. Historically, however, the centralized mainframe computing systems of IBM and others provided a unique architecture to gather processing data. Most examples

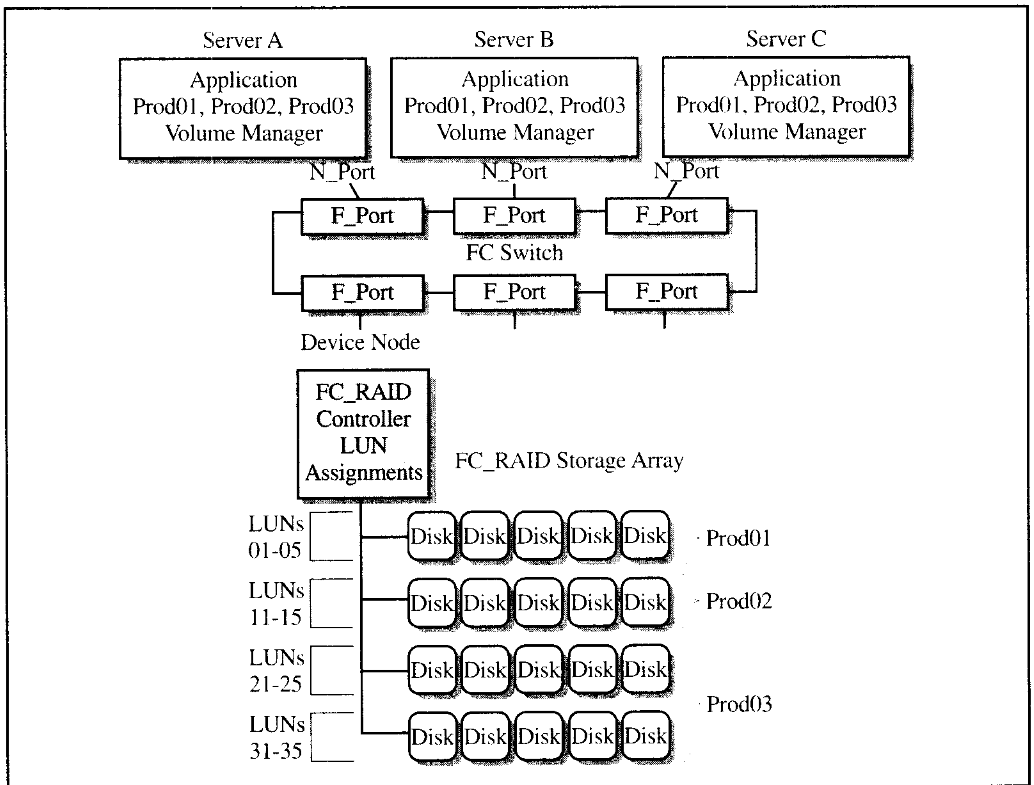

Figure 24-2. *Volume management working with storage arrays*

in this area point to the IBM System Management Facility (SMF) and Resource Management Facility (RMF) functions as models for what should be expected within the open systems area. Overall, this is not the case, and taken out of its centralized and proprietary context, the SMF and RMF models for historical information gathering do not fit well into the area of client/server and distributed computing.

The first area not covered by the IBM model is the collection and support of distributed computing configurations that have become commonplace within client/server installations. This is exacerbated by the open architectures of UNIX systems and lack of any standards in processing nomenclature. Consequently, there are multiple ways of viewing a process within the UNIX variants of the UNIX open systems model. This permeates into the storage space, as the nomenclature and operation of UNIX storage models are, by design, very flexible and open to vendor interpretation. So gathering data about a SUN configuration is a different exercise than gathering processing information about an AIX environment. If we throw in the increased open area of storage, vendor support for each of these environments makes the common historical database such as IBM's SMF and RMF impossible.

The second area that makes the centralized proprietary environment different is the distributed nature of processing. As servers became specialized, the ability to provide a historical record in context with the application became problematic at best. Consider a database transaction that executes partly on application servers, where application logic creates and submits a database query, which is executed on another server. The historical log of the application processing becomes difficult to track as it migrates from one server to another. With the distributed processing activity, the combined activity information transverses from one operating system to another with the related resource utilization. This example points out the problematic recording of database transactions within the context of the application.

We will integrate the complexities of a SAN into the collection discussion and provide yet another source of processing and resource information into the capacity planning equation. In terms of a SAN, this becomes the fabric and related micro-kernel operations in moving the data within the storage network. If we integrate the NAS environment, we integrate yet two more sources of processing: the NAS micro-kernel processing and the operating systems that run the network fabric. The collection of data for both SAN and NAS adds additional data collection points that are again distributed in nature.

The point is that the collection of data within the distributed environment continues to be problematic. Although there are some methods that attempt to deal with these issues, the need to develop an alternative to the historical collection of data becomes evident. Currently, only a limited number of software tools are effectively addressing this area using emerging standards that not all vendors have accepted. This exacerbates the situation for storage analysis within storage networks, given the difficulty in collecting valid storage activity information.

However, there are two basic initiatives to consider when collecting historical storage activity information. The first is the industry initiative of the common information model (CIM) that is an object-oriented database for collecting system processing details and device status. CIM is a Microsoft-initiated standard that has achieved some level of acceptance although its complexity continues to make it impractical for storage network usage. The other item is the usage of network-oriented tools which many storage-networking vendors include with their products: the management information base (MIB). These provide the quickest way of collecting information in-band for the SAN and within the NAS micro-kernel (although this isn't used as often in the NAS micro-kernel levels).

The Common Information Model and Storage Networks

The common information model (CIM) is an object specification that provides a uniform way of describing a computer and its components. The specification is used as a standard to write applications that access the objects as described within the model. The CIM initiative first started as a Microsoft object-oriented structure used mainly by the operating system to build a repository of general and specific information about the computer it was running on.

The objective of this functionality was to give hardware companies the ability to access their components within a network. Used for support and problem management, vendors quickly contributed to the Microsoft-led initiative with specifications for general computer components, concentrating chiefly on internal components for processors, adapter cards, and memory structures. However, it was also Microsoft's objective to evaluate a configuration in order to determine the requirements needed for particular operating system functions, applications packages, and licensing verification. Oddly enough, the initial releases of the CIM specification did not reflect any of the major peripheral and processing software components necessary to complete a computer configuration.

The Microsoft CIM initiative was soon passed to the Distributed Desktop Management Task Force, DMTF, a standards-based body that manages similar initiatives for distributed PCs. DMTF had already recognized the CIM initiative and included it as part of the initial releases of DMTF standards. The move to integrate CIM as a DMTF standard came about the time another consortium standard initiative was passed into the DMTF—this was WBEM, or web-based enterprise management.

The WBEM initiative began the development of particular standards for enterprise management of computers over the Web. As Microsoft was one of the original members of the WBEM consortium, one of the standards initiatives it began with was the CIM standard. So any vendor wishing to develop products that could be managed uniformly, regardless of vendor association, and managed over the Web, must meet the CIM and WBEM standards.

Consequently, specifications for general storage devices (say, IDE and SCSI drives) and offline media, such as tape and optical peripherals, were not added until later. The component descriptions for Storage Area Networks are very new and have only been added as of the writing of this book. However, given vendor acceptance and cooperation, CIM and WBEM do provide a uniform description for storage devices and storage network components. As these specifications are incorporated within vendor products, the data center begins to see a consistent view of the storage infrastructure.

The important points here are the eventual scope of the CIM specifications and the use of CIM as a standard. As a standard, that means sufficient vendors must integrate their product offerings into the CIM standard before the concept becomes relevant to the data center. As a meaningful solution, both vendors that develop management products as well as vendors that develop storage networking products must adhere to the CIM standard for this to be productive. In other words, it would be difficult if you have vendor A's disk products that conformed to the CIM standard, and vendor B's disk products that did not but conform to their own specifications of storage devices and activities. The same can be said for vendors who provide products for performance management and capacity planning—the standard must support all instances of CIM implementation at the storage hardware and software level. In addition, the CIM specification itself must have sufficient detail to provide value.

SNMP and MIBs

MIBs, on the other hand, have come through the standards process and been accepted as a valid solution for collecting and describing information about—get ready for this—networks. Yes, the MIB specification came from the network environment and continues to be used today as a fundamental element of network management activities.

Management information bases (MIBs) are very complex file-oriented databases that describe a network and its components and act as a repository for activity information that occurs within the network. The MIBs were the database for a distributed protocol that is used to access remote networks for management purposes. This protocol is known as the Simple Network Management Protocol, or SNMP. Anyone with networking experience should be quite familiar with this concept and protocol.

SNMP and their related MIBs create a way of collecting information for inclusion in performance, problem, and capacity management. However, SNMP and their related MIBs are complex systems that require specific programming to derive value. Many network management products base their functions on SNMP and MIB standards. This proximity of network technologies within the SAN environments prompted the inclusion of MIBs within the SAN switch software.

Leveraged by SAN vendors as an accepted network management tool, the SNMP and MIB combination for SAN management laid the groundwork for today's management repository for SANs. Although not as likely to be accessed or included within the micro-kernel applications, the NAS configurations can also be included in SNMP and MIB solutions. The caveat is the complexities of the solution and the long-term viability of the MIB. The data center must rely on third-party software products that integrate SNMP functions while allocating their own MIB files, or face writing their own SNMP scripts and defining the MIB files themselves. The latter is a complex, time-consuming task within the data center.

Modeling Performance and Capacity Requirements

One objective of the capacity planning activities is to estimate as accurately as possible the future needs of users. This generally can be accomplished from modeling the potential workloads and configurations for adequate fit. In general-purpose capacity planning terms, there are levels of modeling that can fit the needs of most data-center requirements. These solutions range from expensive and specialized software tools that model workloads and potential performance based upon hardware configurations, to vendor benchmarking that generally uses subsets of the actual data-center workloads, and the official performance councils that provide third-party benchmarks of vendor solutions.

The first set is the capacity planning modeling tools. Although these provide the greatest detail, their accuracy depends on the data applied to the model, and the expertise in analyzing the output of the benchmark results. Unfortunately, these tools have yet to be developed for storage networks, but through creative efforts could probably be modeled given sufficient time and money. Even then, the model would be an integrated simulation of workloads that share common storage arrays, or in the case of NAS workloads, that provide additional latencies as remote network drives.

The second set is the vendor benchmarks, although these by their very nature will be suspect given their inability to replicate the specifics of an individual data center. These simulations don't always have the disparate facilities that make up production data centers, and as a result, the benchmark may be skewed toward the vendor's solution. Wouldn't that be a surprise? However, vendor's benchmarks provide valuable insight into understanding the potential capacity and performance of an expensive storage infrastructure installation. The additional aspect is that many first-tier vendors have user benchmark centers where they test potential customer solutions as well as conduct their own interoperability testing.

The third set is the third-party benchmarks by non-profit corporations that sponsor testing and performance benchmarks of real-life configurations. These companies are likened to the insurance safety councils that perform crash tests. The performance councils take off-the-shelf equipment from vendors and build a real-life configuration in order to run a simulated workload based upon end-user applications, such as OLTP and data warehouse transactions. In other words, they test out the configuration in real-life scenarios so as to validate all the factors a data center would consider when purchasing the configuration. Two are relevant to the storage industry: the Transaction Processing Performance Council (TPC) and the Storage Performance Council (SPC).

The TPC provides benchmark testing of computer configurations using standard transactional sets. These benchmarks execute transactions that characterize database queries which simulate everything from simple queries to complex data warehouse queries that access multiple databases. The tests are run on vendor-supplied hardware and software configurations that range from homogenous hardware systems to heterogeneous software operating environments and database systems. The test results are generally published and available for purchase through the council. This allows data centers to monitor different levels of potential configurations at arm's length while obtaining information about potential cost-to-operating environment requirements. This provides an evaluation of storage from an integrated view, as storage configurations become part of the system's overall configuration.

The SPC is specific to storage and is the new kid on the block when it comes to evaluating vendor storage configurations. This is the most specific and productive modeling available to date for storage networking and capacity modeling. Their job is to be the insurance safety council for the storage industry and protect the data center from products that continue to be problematic in real-life implementations, while providing an effective feedback mechanism for vendors who strive for better goods.

The SPC-specific objectives are meant to provide both the data center and systems integrators with an accurate database of performance and price/performance results spanning manufacturers, configurations, and products. They also use these experiences to build tools that help data centers analyze and effectively configure storage networks.

They do this through a series of configuration requirements, performance metrics, and tests. The services can analyze small subsets of storage, from JBOD and RAID storage arrays to large-scale SAN configurations. However, all configurations must meet the following criteria prior to testing:

- **Data Persistence** Storage used in an SPC test must demonstrate the ability to preserve data without corruption or loss. Equipment sponsors are required to complete audited tests that verify this capability.
- **Sustainability** A benchmark configuration must easily demonstrate that results can be consistently maintained over long periods of time as would be expected in system environments with demanding long-term I/O request throughput requirements.
- **Equal Access to Host Systems** All host systems used to impose benchmark-related I/O load on the tested storage configuration must have equal access to all storage resources.
- **Support for General Purpose Applications** SPC benchmarks provide objective and verifiable performance data. Specifically prohibited are benchmark systems whose primary purpose is the performance optimization of the SPC benchmark results without corresponding applicability to real-world applications and environments.

Vendors who submit their products to these benchmarks must have their systems available to ship to customers within 60 days of reporting the SPC benchmark tests.

Probably the most valuable aspect of the SPC benchmarks is the actual test. The SPC has developed two environments that depict many of the workload demands we have previously discussed (see Chapter 17). The following describes two test scenarios that are run.

- **SPC1 IOPS (I/Os per Second) Metric** An environment composed of application systems that have many users and simultaneous application transactions which can saturate the total I/O operations capacity of a storage subsystem. An OLTP application model makes up the benchmark where the success of the system rests on the ability of the storage system to process large numbers of I/O requests while maintaining acceptable response times to the end users.
- **SPC1-LRT (Least Response Time) Metric** This environment depicts a batch type of operations where applications are dependent on elapsed time requirements to complete. These applications provide multiple I/O requests, which are often serial in nature—in other words, they must complete in a predefined order. The success of the storage system in these processing environments is dependent on its ability to minimize the response time for each I/O request and thereby limit the elapsed time necessary.

The SPC carefully audits and validates the results of benchmarks. Configurations and testing criteria is audited and validated either onsite or remotely through an audit protocol. This serves to provide the vendor with audit certification that the tests and configurations meet the SPC standards and testing criteria. A peer review is conducted upon the completion of benchmark results. Results are considered validated and become official upon the completion of the 60-day peer review process if no compliance challenges have been brought forward. Official results are available to SPC members on their web site and open to certain publication rights.

Implementing the Plan

Taking into account a comprehensive analysis of the existing or proposed storage infrastructure, implementation becomes the critical factor. A big factor in implementation is the acquisition of upgrade hardware and software or a completely new solution. Coupled with acquisition of the storage networking components is the scheduled installation, testing, and transfer of production storage data and workload I/O processing. This is followed by the initialization of the systems management cycle of activities.

Acquisition of storage capacity is an exercise in commodity purchases driven in most data centers by price. The best price per MB usually gets the business and drives storage vendors to differentiate their offerings in the firmware and software solutions that become part of the system. However, the storage networking business has turned this upside-down, with the necessity of new networking devices and the enhanced storage arrays and tape systems that must be FC-enabled. This has placed price in a more balanced perspective with other major purchase factors, such as reliability, service, and quality.

Consequently, the following guidelines can be used to assist in data-center acquisition strategies when purchasing storage networking solutions.

- **Competitive Bids** It's important to gather at least three competitive bids for the configuration you are seeking. Ensure that the requirements for storage networking, be it SAN or multiple NAS solutions, are available to the vendors you have decided to work with. Be advised that "total solution" offerings that are provided through storage vendors and systems integrators allow them an additional revenue source.
- **OEM Knowledge** Most storage networking vendors, especially larger system vendors, provide their solutions as a composite of external components supplied by the third-party companies. These companies provide their component, be it hardware or software, as an Original Equipment Manufacturer (OEM) supplier, and place their name on the device. This is commonplace in storage networking with most SAN solutions being made up of OEM suppliers. This is not a bad thing, but it's important to compare apples to apples when considering

competitive bids. In other words, don't pay extra for the same FC equipment, such as switches, HBAs, or routers that is available through a competitive solution.

- **Storage vs. the Network** Try to separate the storage part from the network part. This is especially helpful in SAN acquisition, given that the network portion is likely to expand at a different rate than the actual storage arrays. The storage capacity plan should provide for incremental upgrades through the planning cycle, generally a year in length. During this time, storage array acquisition can be part of the competitive bidding process, but should be negotiated with as little bid lock-in as possible. In other words, have the vendors consider bidding on the entire plan (for example, a complete year of requirements) to obtain a better price, service, and vendor commitment.

Design and Configuration

Many of the examples shown in the book depict storage networks supporting three common types of workloads: OLTP, Web Internet Based, and Data Warehouse. More details on these I/O workloads can be found throughout the book with additional guidelines on identification and estimating in Chapter 17. Typically, SAN configurations are comprised of combinations of 8-, 16-, and 32-port FC switches, with disk arrays commensurate with storage capacities that have been estimated with workloads. Within NAS installations, typical configurations can encompass both departmental solutions and support within the data center for more sophisticated applications. Another important point to consider is the external factors that influence the installation. These are Ethernet network modifications within NAS installations and should be considered prior to installation. SAN configurations are subject to the inclusion of intersystem-link ports (ISLs) and an integrated FC-SCSI bridge into a tape library.

Test Installation

It is important to define, configure, and install a permanent test installation environment. Putting a small configuration in place provides essential first-case experiences in the configuration and operation of both SAN and NAS installations. This also provides a test bed for testing future software and hardware upgrades while enabling an application testing facility.

Use the test installation to initiate a pseudo-management practice. Management becomes a very challenging activity, especially when operating the new devices and complexities in the SAN. It also is the most rapidly evolving practice, with constant change occurring in software tools and accepted practices. Additional discussion of storage network management topics can be found in Part VI.

Production Installation

Develop a production turnover activity where a formal change window is established. In many cases, this may need to be integrated into existing change management activity within the data center. Key among these is tracking the changes made to all components of the storage network (see Chapter 23). It becomes particularly troublesome if you formalize changes to the switch configurations and not to upgrades of critical components such as HBAs, routers, and attached storage devices.

An important aspect of product installation is establishing a backout practice. Because the storage network is an infrastructure in and of itself, the reliability can be problematic in the beginning, as with any new technology. However, being able to back out quickly and return the production environment to an existing state saves valuable time as you move into a storage networking environment.

Closing the Loop

Here, we come full circle. Back to the place we started in this part of the book—that is, a sound and well-thought-out capacity plan for storage networking. The benefits are evident in the increased availability, reliability, and success of the storage networking infrastructure. In addition to setting the basis for a manageable set of service levels, the capacity plan supports the external systems and applications departments in the data center. More important is the credibility it establishes for that support.

The rationale that storage networking is a new technology and does not lend itself to traditional or summary planning for capacity is incorrect. As we have demonstrated, the process of estimating resources and configurations can be done. However, it is an iterative process and one that becomes more accurate as mistakes and related information arise. As new information and tools become available, the accuracy and time required to plan for effective storage networking configurations will increase. Key among these will be the increased information collection methods and the ability to provide effective statistical predictions without the collection of months and months of detailed activity information.

This must be balanced against the fast-moving storage networking industry where new technologies become available almost every calendar quarter. The difficulties behind planning without the knowledge of technology advancements will move both the vendor community and the data center to reach some level of product planning communications that is beneficial to both parties. As discussed in Chapter 20, new innovations continue to affect the storage networking solution. Among these is the coalescing of SAN and NAS architectures and the advancement of traditional computer bus technologies into switched fabrics.

Storage networking has become the cornerstone of new applications within the data center and continues to deliver the scalable performance necessary for the future. However, these new storage functions must be applied in an organized and planned fashion. Capacity planning for storage networking infrastructures may form the foundation for the next level of storage management, culminating in a true paradigm shift within the data center.