Chapter 22

Managing Availability

Availability is driven by capacity and performance. Without consideration for matching the capacity of the storage size and configuration to the workload, availability becomes a guessing game. Consequently, there must be a sound basis for the existing storage capacities and configurations. However, taking reality into consideration, using the most overlooked work axiom—*do it right the first time*—the existence of a sound capacity plan for an initial configuration may not be the case. On the other hand, storage networking, SAN more than NAS, is a new solution and a learning curve always begins with mistakes. Therefore, this chapter begins with the assumption that mistakes have already happened, can easily happen, and in the best case, can be avoided.

The premise of availability for storage networks is the application of the five systems management disciplines in an iterative process. These will be discussed as they occur in a real-time continuum, but we'll also view them in a cyclic manner. The following illustration offers a peek into the activities that make up the cycle. As we can see, a capacity model and implementation may set the basis for the start of the cycle, but may not have much to do with the reality of daily data-center operations. Performance management is influenced by both configuration and change activities, and given its investigative characteristics, forms a logical beginning to availability.

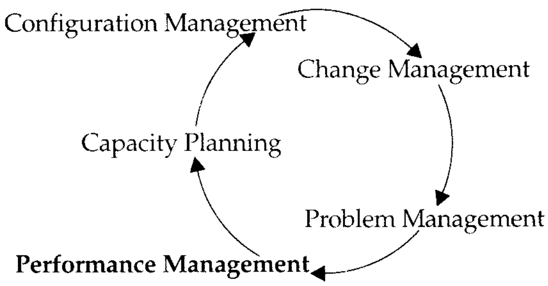

The performance management activities are supported by external events that are tracked through problem and application management systems. All of these activities develop a synergistic relationship that cycles back to influence the capacity model. So you see, it is, and continues to be, an iterative process in providing available systems.

This chapter discusses these interrelationships as they are related to the SAN and NAS environments. By applying activities and concepts from systems and network management practices, the availability of storage networks characterizes their fundamental difference compared to existing direct-attached storage models. In addition, these discussions begin to demonstrate the immaturity and challenges still evident within this area of storage infrastructures. Focusing on both fundamentals of systems, and storage management, the scope of these discussions illuminate key issues and offer insight into providing availability in storage networking environments.

Availability Metrics

Providing a measurement of services within the storage network requires some type of metric. The most encompassing is the service level. Service levels exist on several levels. There are specific agreements with end users that define the level of service the IT department provides. There are internal service levels that are managed to provide a working measurement of systems and networking performance. Indeed, working with existing service levels for storage, if they exist, may be either too specific or fail to have sufficient depth to cover performance beyond pure capacity measurements. The key is the establishment of meaningful measurements for the SAN infrastructure that encompasses both storage systems and network specifications.

Establishing service levels for storage network configurations requires the integration of dependent system functions, related network activities, and specific storage elements as they come together to form the storage network infrastructure. First and foremost in the development of these metrics is the association of the workloads. A discussion of estimating workload I/Os can be found in Chapter 17. This is an important exercise because it forces discipline regarding establishing the user's side of the equation and assists in integrating storage configuration service levels into supported applications.

As discussed, two forms of service levels should be developed: an external articulation of services to the supporting business units and end users, and an internal service level that supports both the data center infrastructure and your application development colleagues (DBAs, other systems administrators, and business continuity teams). The external service to the business units is explained in the following section.

External Service Level for Storage Networks

The following can provide a data availability framework for working with external service levels.

Data Availability

This component defines which data is referred to with the service and when it will be available. In addition will be a understanding as to when the data will be unavailable. All data becomes unavailable at some point because of the need to back up files and volumes periodically, and because of the occasional system upgrade that makes all data unavailable while hardware, operating system, or application software is upgraded. These are typical of general service levels and should be included within those commitments. For storage service levels, however, the data center should note that particular availability requirements will be subject to additional or similar instances where the hardware, switch OS, or a management application will be upgraded. These may require the user data to be unavailable during those times.

Data availability is extremely important to the data center. The ability of meeting external service levels will center upon IT activities to define the service levels along with supporting the service levels.

- **Uptime** This metric is usually expressed as a percentage of the data availability period that the data center has committed to. This is driven by the need to keep data available for particular applications, user transactions, or general-purpose access (as with network file systems), and reflects the sensitivity of the business application to the company's operations. Examples are expressed in terms of 99.999, the five nines to distinguish that the data is committed to be available for that percentage of the overall time period it has been committed to. In other words, if an OLTP application requires a 99.999 uptime during the transaction period of 6 A.M. to 6 P.M. PST to support the operations of both coasts, then the percentage is applied to that 12-hour period. Therefore, in this example, the OLTP application requires the data be available 719.9928 minutes of that 12-hour (720 minutes) time period—effectively rendering this a 100-percent uptime for the data.
- **Scheduled Downtime** This metric defines the known time required to perform periodic maintenance to the data. Knowing this up front provides additional credibility to the storage infrastructure and reflects the time the data is unavailable to end users and applications. The important distinction is that the same data used in the transaction system very likely will be the same data that is accessed in batch mode later that night. Although transparent to the end users, this can be critically important to internal service levels.

Note *It's important to note the distinction between data availability uptime and response time. The main purpose of the storage infrastructure is to keep the data available for the required applications. Given that actual application response time is made up of a series of components that include the operating system, database, and network, the ability to provide metrics beyond the availability of users is beyond the storage infrastructure's scope. However, the commitment to I/O performance must be reflected in the internal service levels leveraged as a component of the overall response time.*

Data Services

This component of the service level defines the services that the storage administration offers.

Although the services listed next are traditionally accepted services of the data center, the storage infrastructure will accept increasing responsibility in meeting service levels surrounding these items.

- **Backup Service** Operations to copy specific sets or groups of data for recovery purposes due to data corruption or service interruptions.

- **Recovery Service** The corresponding recovery operation that restores data compromised due to data corruption or service interruption.
- **Replication Services** Services provided by the storage infrastructure that send copies of user data to other processing infrastructures throughout a company. Generally used to copy data to remote offices for local processing requirements.
- **Archival Services** Services providing periodic copies of data for necessary legal, local policy, or governmental archival purposes. Used in conjunction with the backup service, they provide a duality of service.

Disaster Recovery Commitments

As part of the business continuity plan, these services and service levels offer data availability during the execution of the disaster plan.

Consider the key storage items listed next when developing service levels for business continuity and disaster recovery plans.

- **Data Recoverability Matrix** It's unlikely that all data will be covered under a disaster site recovery plan. Therefore, a matrix, list, or report should be available regarding what data is available at any point in time prior to executing the disaster site recovery plan. The service levels and services here are to some degree a subset of services available during normal operations.
- **Data Uptime** The same metric as the normal operating data uptime percentage and predefined processing time period.
- **Data Scheduled DownTime** The same metric as the normal operating percentage, but using a predefined processing time period.
- **Restoration of Services** A new level of service driven by the schedule for total restoration of services for a disaster site recovery plan.

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 unexpected circumstances do happen and should they occur, they should be resolved with mutual understanding and changes in the existing plan. This can be critical given the volatility of SAN installations and NAS data-volume requirements.

The following are a few of the unexpected circumstances that are worth discussing with end-user clients. These can form the basis for a non-aggression pact that will benefit both parties in managing to the agreed upon service levels.

- **Unexpected Requirements** The most common set of unexpected circumstances is the unforeseen application that requires support but which hasn't 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 and 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 the initial SAN design proves insufficient to handle the workload. The requirement to retrofit the SAN configuration with enhanced components means additional cost and disruption. This can at least be addressed with the understanding that new technology installations are available.
- **Mid-term Corrections** It is likely that any enterprise storage installation will experience either one or both of the preceding conditions. Consequently, it is extremely important to build into the user agreements an ability to provide mid-term corrections that are an evaluation of the current services and corrections.

Internal Service Levels for Storage Networks

Within the storage infrastructures, and especially the storage networking areas, the support is both to end users as well as other individuals in the data center. Therefore, there will be an internal set of service levels that support the other infrastructures within the data center. These at a macro level are the systems' organizations, of which storage may be a component, the support of systems administrators, and from which web masters, systems programmers, and system-level database administrators (DBAs) will be critical. On the applications side, storage remains an integral part of applications programmers, analysts, and maintenance programmers' requirements. Certainly, we can't forget the network administrators, help desk, and network maintenance people.

The following is a more detailed discussion on best practices for negotiating and determining service levels for internal IT co-workers and management.

- **Storage Capacity** The most common requirement for the data center is the raw storage capacity necessary for applications, support processing, and database support, just to mention a few. A service level here will prove very productive in staying ahead of user requirements and system upgrades that require additional raw storage capacity. This is critical with SAN and NAS configurations, given that each has attributes that make upgrades more support-intensive. With the manipulation of switch configurations, zoning, and LUN management upgrades, storage capacity in the SAN environment is not a "snap on" activity. Although NAS provides a much easier way of increasing storage, the bundling of the storage device may provide more storage than required and external network performance effects may be more intense and challenging than they first appear.
- **Data Availability** This service level is similar to the end-user service level, and in many ways forms the foundation for the application service level with the same end user. Over and above the commitments for production data,

which should be addressed with the supporting applications personnel, is the need for data to be available for testing, quality assurance functions, and code development.

- **Data Services** Again, similar to the end-user services specified in the external functions, are the services provided to assure backup/recovery and data archiving. These services include operating system copies, code archives, and backups for network configuration data to name a few. Relating these to the SAN infrastructure initiates an increase in the complexities of recovery operations. Given these installations are new and require additional procedural recovery operations to be developed, a learning curve is expected.
- **Storage Reporting** Key to internal constituencies, this provides the basis for monitoring data usage and activity within the storage network configurations. It sets the stage for effective housekeeping activities by establishing either new or preexisting quotas, archival policies, and ownership tracking. In addition to these administrative tasks, it establishes a tracking mechanism to help fine-tune the storage configuration. This is key to the SAN environment where volume allocation, physical data placement, and systems/application access are controlled through the storage support personnel and your configuration.

Implementing the Plan

Now comes the fun part: the implementation of the storage network configuration. It's time to develop or redevelop the SAN or NAS configurations to suit your particular service levels. The configuration options, driven by service levels and both internal and external requirements, demand a particular level of availability. These are generally categorized as on-demand availability, highly available systems, or fault-tolerant systems. In the end, sensitivity to the end user is an external call and end users obviously have final say in determining their value. These categories are discussed next.

- **On-Demand Availability** These systems support 80 to 95 percent availability for defined processing periods. They generally support transactional systems that do not significantly impact the daily operations of the company if they are down for a reasonable time. Most of these systems are data warehouse systems that are prone to outages given the level of data and their reliance on background processing. Others are those that provide purely informational services to internal users of a company, such as intranet systems, inventory, and educational systems. Figure 22-1 shows a typical cascading configuration supporting a data warehouse.
- **Highly Available** These systems require 99.999 percent uptime, and given the right resources they can achieve this. They are characterized by applications that have a significant impact on the company's business if they are down. This includes systems such as OLTP financial and banking applications (for obvious

reasons), retail and point-of-sale transactional systems, and customer relationship systems that allow customers to place orders, as illustrated in Figure 22-2.

- **Fault Tolerant** These systems cannot afford any downtime and must be available 24/7. This includes transportation systems such as those supporting air, emergency, mass transit, medical, and emergency 911 systems. For obvious reasons, these need to be supported by fully redundant systems that can provide full failover capability so operations remain uninterrupted. Figure 22-3 depicts a redundancy mesh SAN configuration supporting an emergency medical response application.

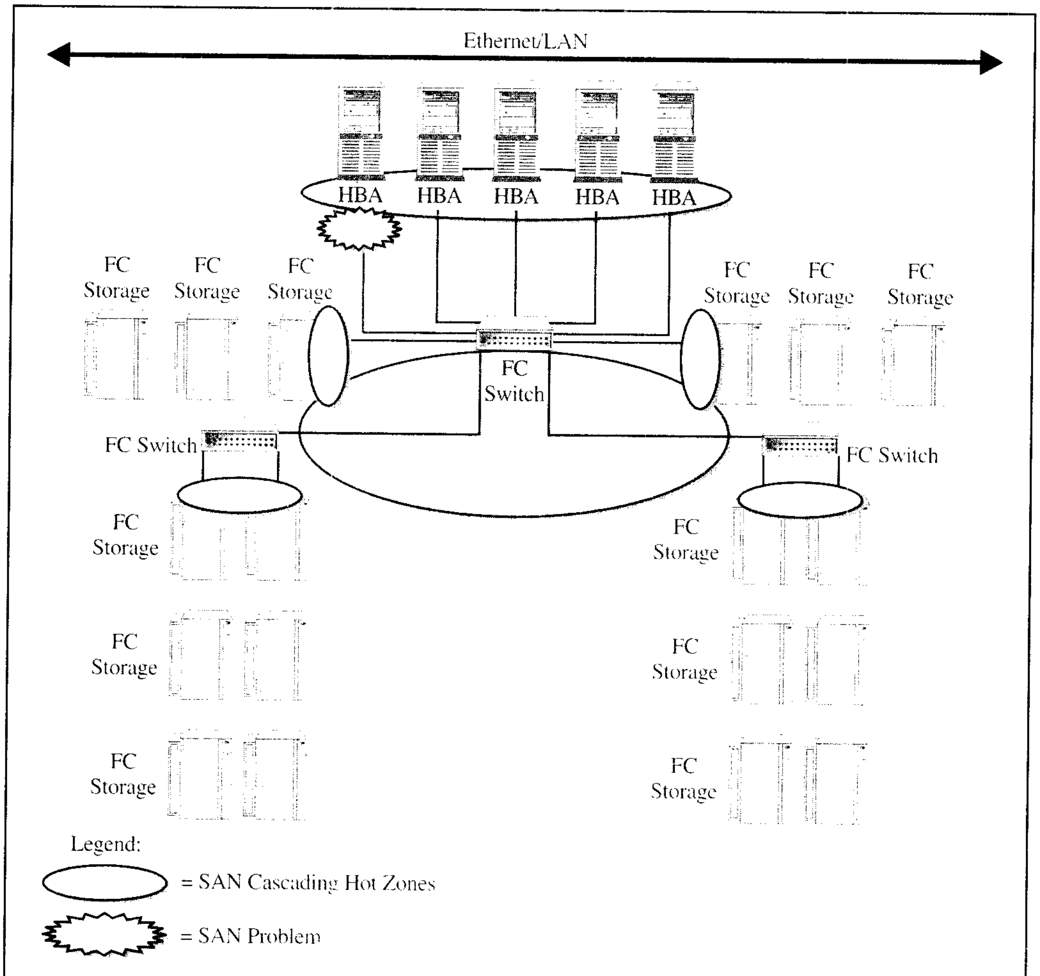

Figure 22-1. *A cascading SAN configuration supporting on-demand availability*

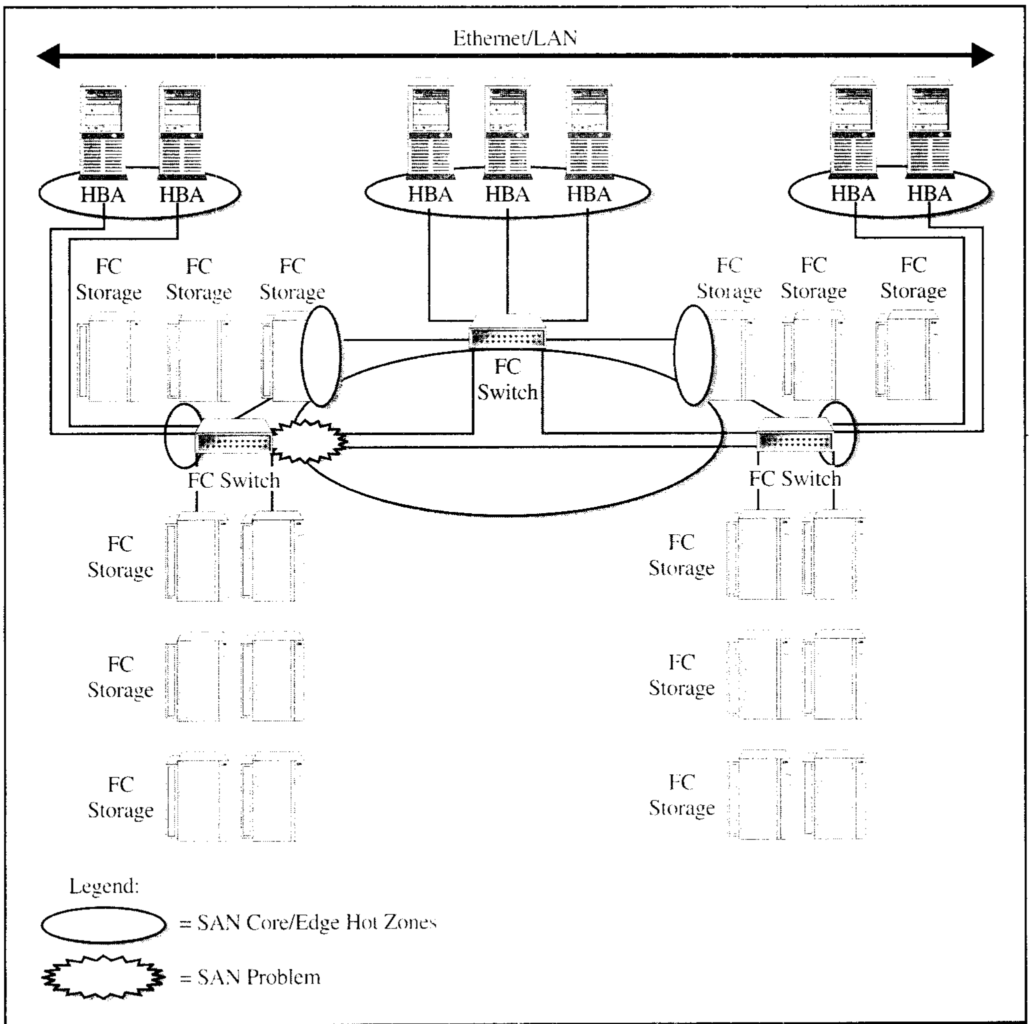

Figure 22-2. *A core/edge configuration supporting a highly available CRM application*

If we contrast these availability categories to SAN and NAS configurations, we find some simple guidelines can be applied to orient them toward appropriate configurations.

On-Demand Availability—NAS and SAN

NAS configurations for on-demand availability can be contained with entry- and workgroup-sized NAS devices. Availability levels in the 80 to 90 percent uptime range are generally acceptable in supporting shared files and network file services for PCs.

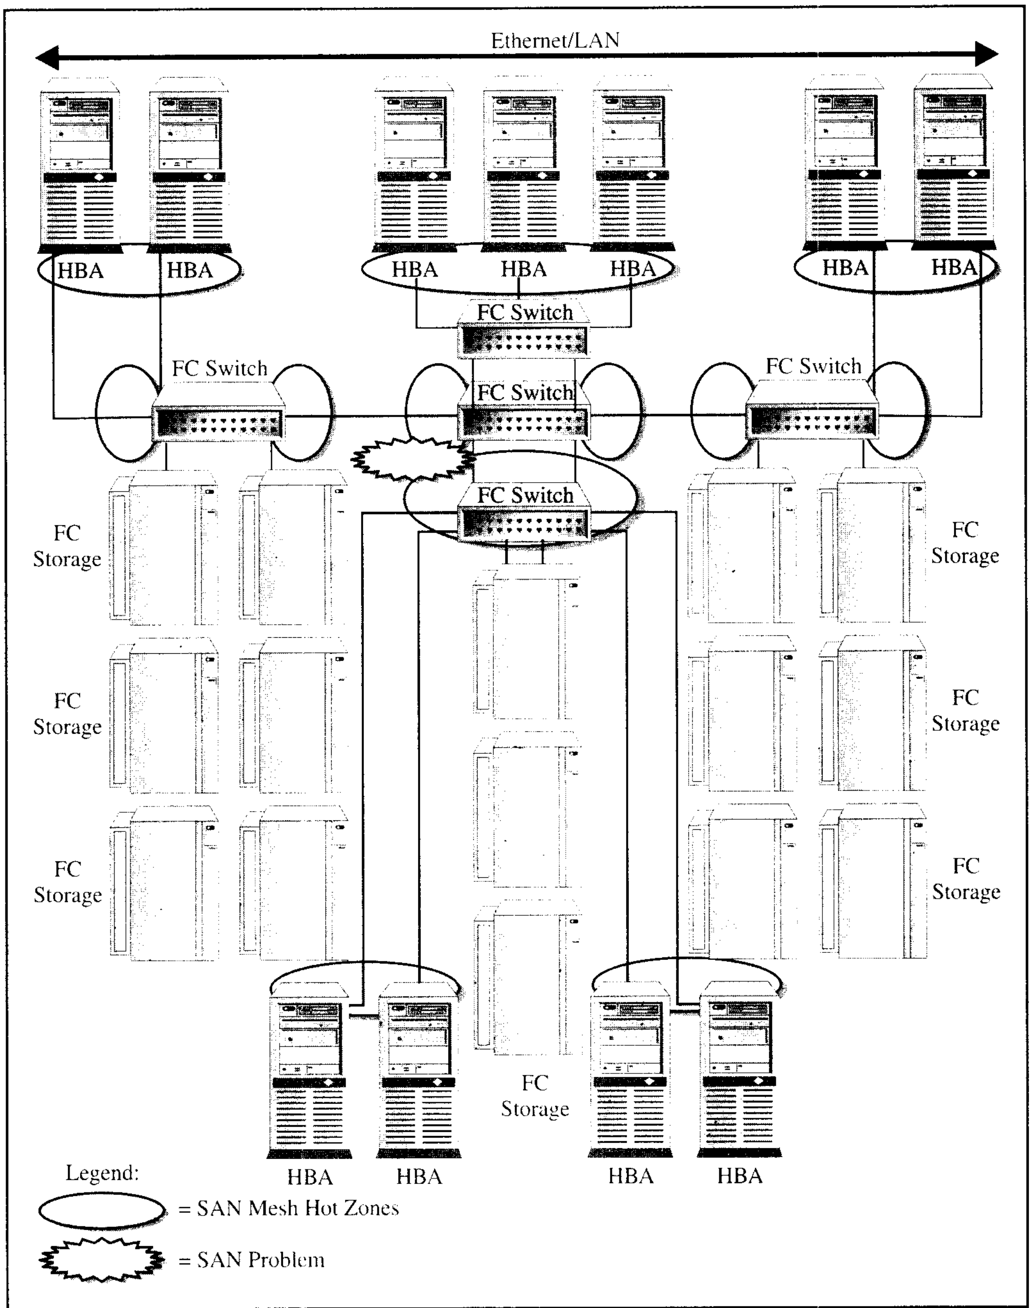

Figure 22-3. *A redundant mesh SAN configuration supporting a 24/7 application*

Given the reliability of the NAS devices performing this type of work, the storage availability will likely be close to the five 9s; however, the variable in keeping these service levels is the external effects of application servers and the local network. Given that scenario, if the external effects are managed within the internal service levels through agreements with your colleagues in the systems and network areas, and external agreements with end users (see Chapter 24) make these effects known, the NAS entry and workgroup devices will perform well.

SAN configurations supporting these areas can be simple or extended cascading configurations where the level of path redundancy is not required, and the performance factors are more oriented toward availability of the data. This will largely be a matter of the reliability of the FC storage devices rather than the switch configuration; however, with today's RAID array functions and their support of larger data segments, the reliability of these devices will be quite high. Because of this, data uptime can exceed 90 percent.

Highly Available Systems—SAN maybe NAS

For systems that require the five 9s, the more robust the underlying devices and the more flexible the configuration, the better the solution. As these systems will likely support the use of RDBMSs for data organizational models, the use of NAS becomes somewhat problematic. In terms of SANs, this means an entry point at which the robust nature of its storage devices and the flexibility of its switched fabric begins to show real value. However, the trade-off is increased complexity, as support for these availability models requires a core/edge configuration to ensure performance that includes consideration for interswitch duplicate paths that are necessary for redundancy. In some cases, the entire SAN configuration may need to be duplicated within a cluster operation, whereby a standby SAN configuration holds the failover system. In the end, this is very expensive and not recommended unless absolutely necessary.

Fault-Tolerant Systems—SAN only

Systems requiring full redundancy may either be highly proprietary systems that provide multiple levels of failover (such as 4- to 8-way clustering systems with OS extensions to synchronize the failover functions), or systems with a single failover used solely as a standby. Other systems may be fully fault tolerant in a tightly coupled system of MPP or high-end SMP architectures (see Chapters 7 and 8 for additional information on MPP and SMP configurations). Regardless of their configuration, these systems require storage. Many of these systems, being UNIX-based, can participate in SAN configurations leveraging the ability to access either shared storage at the hardware level, more sophisticated (and proprietary) shared data systems at the software level, or a combination of both. This brings the value of SAN configurations further into focus.

A word about NAS when used in fault-tolerant environments. It's not true that NAS solutions cannot be used in these environments; there are several caveats that must be considered. Among these are the uses of an RDBMS within the application, and subsequently in housing the database system and user data. If that's the case, NAS is more likely to be problematic given it has limited abilities with the relational data organizational model (see Chapters 9 and 11 on NAS architectures and NAS Software, respectively). In addition is the interconnection using a TCP/IP-based network with a speed fast enough to support failover operations providing a mirrored data environment. If this requires the installation of a short hop special 10Gbe switch and cabling, the cost may or may not be worth the effort. Lastly will be the flexibility affording the switch's network configurations. Each switch should have accommodations for redundant paths to handle additional as well as alternative paths for data communications.

Storage Availability—Considering RAID

RAID storage configurations are available for both SAN and NAS configurations. Certainly from an availability perspective, RAID has become a requirement for most data centers. Given its implementation, reliability, and cost, it has become a defacto standard for applications. Thus, the question becomes: what level of RAID to use for particular applications—or in terms of this book, what level to use for the I/O workload.

Here, RAID is used in basically two configurations: level 5 and level 1. Other RAID levels are used for specific application support like RAID level 10, also referred to as 0 +1, where files are duplicated with higher performance. However, this comes at the sacrifice of non-stop reliability given there is there is no parity information calculated or stored in this configuration. Other types, such as RAID level 4, are bundled into some NAS vendor solutions.

Considerations for RAID and its associated applications and availability include the following:

- **RAID Level 5** This level of RAID provides data striping on each device within the storage array. The parity information—the information the RAID software or firmware uses to reconstruct the data if a disk is lost—is also striped across the array. This configuration is generally used for high-transactional OLTP applications that require a high-availability service level. The caveat is the write-intensive nature of the application that will increase I/O latency given the increased write operations, as well as the storage array for the multiple writes necessary with the data and the parity information.
- **RAID Level 1** This level of RAID provides data mirroring with parity so the application can continue to run if one of the disk pairs is inoperative. This requires more storage capacity given it's a true duplication of the data, thus allowing a 2 to 1 ratio for data capacity (in other words, one byte of user data requires a duplicate or mirror). Consequently, a 100GB database table that uses RAID level 1 would require 200GB of storage on different disks. Additional data requirements for parity information are not significant and should not be taken into consideration. RAID level 1 is used for applications that require

failover availability, or data-sensitive applications that require data to be available in case of corruption. This comes in handy with applications such as data warehousing, and some OLTP using failover configurations.

Finding the Holes

The most challenging aspect to managing storage networking availability is collecting, reporting, and analyzing the information, while the most important aspect of business application performance is accessing data. The data may be stored on disk, tape, optical, temporary cache, or other device. The location of the data becomes irrelevant; however, if the application cannot get to it, or the paths to the data are congested or critically impacted, the performance of the application will degrade, or worse, become unavailable.

The key to successfully monitoring and managing the performance of business applications is the consistent and proactive management of the critical paths to application data. Although software tools are becoming available, they provide a disparate, incompatible, and inconsistent view of storage information, not to mention storage networking configurations. No single tool provides consistent, proactive management functions that associate business applications with application data. IT management must choose from an assortment of tools that provide only discrete levels of empirical information, ranging from operating system metrics and database metrics, to I/O and disk metrics. IT users bear the burden of correlating these seemingly unrelated sets of information in an attempt to understand the effects of workloads on storage networking resources.

Where to Look

The deficiencies within storage software management tools are compounded by the requirements, costs, and expertise needed to support an increasing set of server platforms, operating systems, and major application subsystems such as relational database management, messaging, and transactional systems. The following points illustrate some of the challenges in managing storage network configurations, as well as the inefficiencies that contribute to business application availability.

- **Correlation Functions among Distributed Components** Storage networks are distributed—meaning events happening on one node device or fabric switch can seriously degrade performance throughout the entire system. The ability to correlate important aspects of performance information as it effects the business application currently must be performed as a set of manual tasks.
- **Proactive Trending** Today's IT managers are expected to effectively drive the bus while monitoring performance through the rearview mirror. Literally all reporting and trending is historical. The information which ultimately reaches the IT user is past tense and provides little value in determining real-time solutions to poorly performing business applications. Consequently, a response to an availability problem with a production configuration may require significantly more time to address.

- **Identification of the "Root Cause" of a Problem** The effects of the conditions stated previously means it is unlikely that the information discovered and reported to the IT user will provide any sort of root cause analysis. This makes the problem management aspect of availability problematic. The information provided to identify and correct the problem would likely address only the symptoms, leading to reoccurring results.

Finding, collecting, and reporting performance information as indicated in the preceding section will be difficult. As such, this should enter into the availability formulae when considering the appropriate configuration and resources. There are multiple sources that IT storage administrators can access to find this information. Given that it remains a manual effort to coalesce the data into something of value, the sources can provide a key to building a historical database of performance and availability information. Particular to the storage networks, these sources are the management information base or MIB (provided within the switch fabric operating system or the NAS RTOS), the hardware activity logs within the specific devices (such as the SCSI enclosures services (SES)), and the OS-dependent activity logs and files that are part of the operating systems attached to the SAN or NAS configurations.

This appears as a loose confederation of information, and it is. However, if one sets up a logical organization of sources, the ability to find, select, and utilize existing tools, and develop internal expertise in utilizing these resources will help a great deal in monitoring availability. The following are some guidelines for both NAS and SAN:

- **NAS RTOS** Essentially everything that is active within the NAS device is available through the NAS operating system. Although vendors have enhanced their offerings for manageability services for monitoring and reporting, they remain proprietary to the vendor and closed to user customization.
- **NAS Configuration** Looking at NAS devices as a total storage configuration requires that external sources of information be identified and accessed. These are multiple and can be defined through network and remote server logs. Again, these require the necessary manual activity to coalesce the data into meaningful information for availability purposes.
- **Fabric Operating System Information** Available within the switch MIBs. (Note that with multiple switches there will be multiple MIBs.) In addition, many switch vendors are enhancing their in-band management utilities and will provide more sophisticated services to access FC MIBs. The same can be said for storage management software vendors that are increasing their efforts to provide in-band software management services that communicate with external products running on the attached servers.
- **HBA Information** Available within the attached node server. However, it's important to note that HBA information coming from system log files must be integrated somehow into a single view. Although this may require additional IT activity, it can and should be accomplished with third-party storage management tools and system collection repositories such as the Common Information Model (CIM).

- **Storage Information** Available within the attached node device through either MIBs, SES utilities, or activity log files generally operating within the RAID control units.
- **OS Information** Available within the attached node server and dependent on the OS (for example, UNIX or Windows). There are multiple sources within UNIX environments to support, find, and select storage performance information. This also includes existing storage management products that support out-of-band management processing. Within Windows environments is a relatively new source: the Common Information Management database. This is an object database that provides activity information for all processes within the server and has become compatible with a large number of storage networking and system vendors.
- **Application Information** Available within the attached servers are multiple sources of information depending on the type of application and its dependent software elements (for example, databases, log, and configuration files). Some of this information is being identified within the CIM database for Windows environments to augment third-party management software suites. For UNIX environments, there are third-party systems management software suites with CIM implementation just beginning. The inherent challenge of these approaches is the correlation of the application information to storage management activity and information.

Note *The problem with any of these strategies, although well intentioned, is the increasing overhead these applications have on switch performance, cost, and flexibility. As pointed out in Chapters 11 and 15, the operating capacities within the NAS RTOS and SAN FC fabric operating systems are quite limited.*

Data Recovery

An advantage of SAN architecture is the leverage of node communications within the fabric to increase availability within the data maintenance processes. Within the data center are multiple maintenance and support applications necessary to maintain platform environments. Of these, none are more basic than backing up files and data for later recovery. The historical problem with these activities is the time lost to copy the data from online media to offline media, using tape media in most cases.

This integral data center practice can be broken into two major activities, each with their own problematic characteristics. First is the process of copying data from disk volumes and writing out the data to a tape volume. Given the disparity of the devices (see Chapter 6 for more on disk and tape devices), a performance problem is inevitable. However, it goes beyond device disparity and is exacerbated by the software architecture of the copy process that has been integrated into most backup and recovery software products. The problem is simple. The traditional, though arcane operation, requires data to be copied from the disk and buffered in memory within the initiating server. The server then issues a write operation for the data in the buffer and the subsequent I/O operation copies the data to the tape media mounted on the tape drive.

This double-write and staging process places a tremendous I/O load on the server executing the copy operation while reserving both the disk device and tape drive during the operation. Figure 22-4 shows how this impacts operations during a typical backup portion of the backup/recovery operation.

The second, and most important, part of this process is the recovery operation. As illustrated in Figure 22-4, the backup is the insurance premium to cover any disruptions and corruptions to the current online data, while the recovery operation is the claim payoff so to speak, when a problem has occurred and data needs to be recovered to both an uncorrupt condition and any previous state.

The recovery operation is different from the copy, even though it appears to be the reverse of the operation; it is far more selective regarding the data that needs to be written back to disks. This requires additional and specific parameters to recover data, such as the specific data to be recovered, from a specific time period, which should be restored to a specific state. The most complex of these operations begins when RDBMSs are involved. This is due to the state condition that needs to be restored during the recovery operation in order to bring the database table to a specific state through the processing of transactional log files.

Enter the Storage Area Network. Figure 22-5 demonstrates the capability of devices within the SAN configuration to communicate with each other, thereby allowing many of the server-based, data-centric maintenance/support applications to be optimized. The tremendous I/O load from typical backup operations can now be offloaded from the initiating server. This requires the data copy functions to be performed from storage network node device to storage network node device (in other words, disk-to-disk,

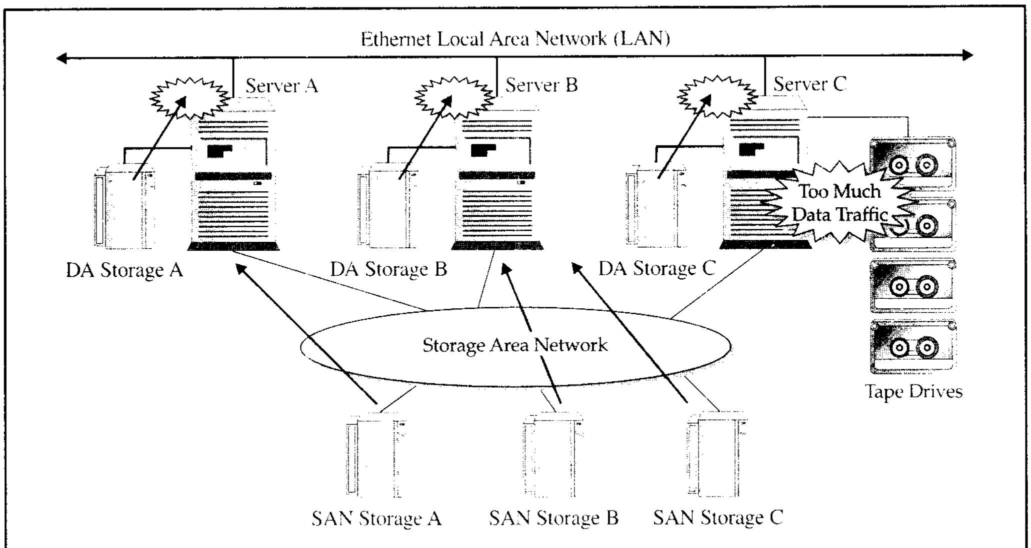

Figure 22-4. *Traditional backup processes and their overhead*

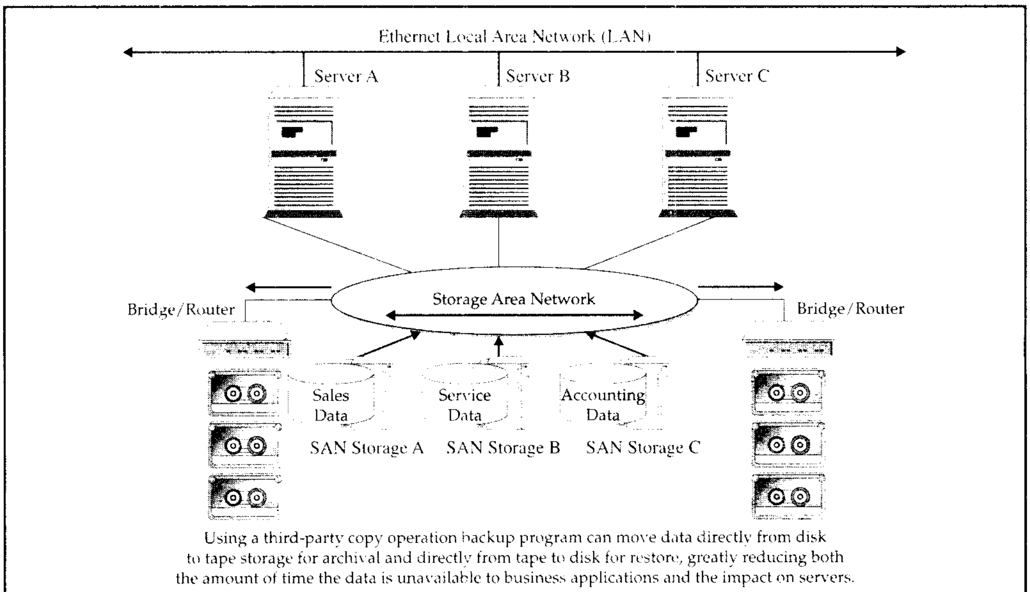

Figure 22-5. *Storage network backup and recovery*

disk-to-optical, disk-to-tape, and so on). Given that the bulk of elapsed time during the traditional backup operation is the double writing of data to the server and then to the backup device, such as tape, this time is optimized through a direct copy operation under the control of the FC fabric working in conjunction with the backup/recovery software that still controls the overall process.

For example, if a typical copy operation used 100 I/O operations from the disk drive/controller to the server path and a subsequent 200 I/O operations to the tape unit, that requires a net 300 I/O operations that the server must perform, not to mention the elapsed time in reserving the disk and tape units. By employing the SAN operation of the direct node communications using the extended copy operation of FC fabric operations, the server I/O can be reduced to a minimum of two to initiate the operation of the fabric in order to directly copy data from the disk to the tape units. The copy operation can thus be performed one time at the speed of FC hardware: 100MB/s with latency for tape bandwidth operations and buffers. Performance depends largely on the specifics of the data, as well as the switch and tape configurations, but suffice to say, the savings will be significant.

Keep in mind that the preceding example was only for the backup portion of the operations. We must also factor in the recovery part. Nevertheless, the savings will be similar to the copy operations—that is, copying the information from the tape to the disk. The operation executes the same way: by the backup/recovery software communicating with the SAN fabric, and the execution of the extended copy command through the fabric. The extenuating circumstances will be the destination of the recovery, the extent of the post processing of transaction logs, and the activity and location of the tape drives that must mount and process the log files.

The significant savings of copying data within a SAN should be taken into context regarding both backup and recovery, because the value in the backup and recovery operation to external service levels is the "R" part, or the recovery. In the end, the time it takes to restore data and application services is key.

Most business applications suffer during this maintenance process because data is unavailable during the time it is being copied. Although necessary for maintaining storage backup procedures and policies, this type of operation (that is, copying data directly from device-to-device) can greatly improve the availability of business applications by reducing the time in which data is unavailable.

Unfortunately, NAS differs in its server-based architecture (albeit a thin server) and its attachment to an Ethernet network. Backup and recovery operations are generally handled by a dedicated server within the subnetwork. NAS vendors have extended their solutions to include SCSI tape attachment. This allows for a self-contained NAS solution which includes its own backup/recovery system. Although it's important for NAS devices in remote locations to have an automated backup process, the capability of NAS to participate within the data center tape library solutions is likely to be a preferred solution for storage administrators.

NAS integration into enterprise-level applications is driving vendors to include data mirroring, snapshot capability, and self-contained backup/recovery operations. These functions have been extended into the NAS hardware and software solutions as well as their integration with FC storage. As this evolves, the ability to participate in device-to-device communication through an extended SAN fabric will make the extended copy operation possible.

Closing the Loop

As available service levels are negotiated, identifying the appropriate and specific information to ensure their availability commitments will become the main responsibility of performance management. Through effective monitoring and reporting of performance factors of the storage network, effective storage network configurations can build from the foundation of information and tools used in this process. With performance management in place, the next phase can be begin to evolve: the effective management of performance anomalies, or problems.

As problems start to be discovered, it will become necessary to ensure that these performance anomalies are accurately identified, tracked, and resolved. Thus begins the problem management phase that illuminates and justifies configuration changes, increases and decreases storage resources, and further modifies and enhances those resources necessary to ensure availability.

Chapter 23 will build on the activities, practices, and challenges illustrated and discussed in performance management and availability.