The Complete Reference

Storage Networks

Part VI

Management—Keeping It Running

Chapter 21

Planning Business Continuity

Computer systems depend largely on their keepers, whether they be embedded systems, the latest blade servers, or the largest super computer. Computer operators and systems programmers have largely given way to increased areas of specialization for administration, maintenance, and servicing. Although many mundane tasks have been automated, more complex configurations of distributed systems have increased the need for system-oriented personnel to monitor and keep the servers running. This is because, no matter how sophisticated they become, computer systems remain electronically controlled components, susceptible to both human and mechanical error.

The basis for these activities revolves around the accepted disciplines of managing systems. These are known as the five systems management disciplines and can be applied to any macro system in operation. These disciplines include the following:

- **Performance Management** The activities associated with the monitoring, tuning, and analysis of a computer system's (or subsystem's) performance.
- **Capacity Management** Analysis activities that drive the planning, design, and implementation of computer resources necessary to maintain established service levels.
- **Configuration Management** Activities that surround the design and implementation of a computer system; this includes ongoing support of system enhancements and upgrades.
- **Change Management** Driven by performance and configuration management activities, these practices monitor and supervise changes that occur to computer systems.
- **Problem Management** Activities that track, analyze, identify, and resolve system problems.

Before you object to the absence of applications management, security management, network management, web management, and others, I should first explain how these topics relate to the subject. The basic disciplines of systems management, as discussed previously, can be applied to any logical definition of a system. Consequently, targeting business applications for management can, and should, have the same disciplines applied to it—in other words, performance, capacity, change, configuration, and problem. By any other name, they usually do. However, particular topics, such as security, add a different dimension to management practices given they are so integrated with the supporting infrastructure. Security, which will be discussed in Chapter 25, should be regarded as a supporting system, like backup and recovery, or database administration. Even here, foundation disciplines can, and should, be applied.

As discussed in Chapter 1, the difference between enterprise applications and programs and utilities that support the systems' infrastructures cannot be overlooked. They form synergistic relationships between the productive and the necessary. Applications drive the business and form the primary components for workloads. However, the performance of computer systems depends on the effective interaction between productive work and overhead. Although many systems programs and utilities are referred to as applications themselves, they often offer productivity enhancements to the systems management activities and therefore to the effectiveness of the IT staff. However, they do nothing for the bottom line. In analyzing basic systems management disciplines, its important to remember that they do not exist purely on their own. Without the bottom line application of computing for some productive processes, the activities that monitor, track, and analyze computer systems would be non-existent.

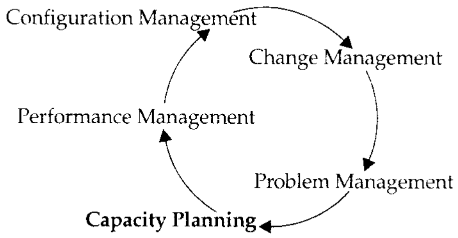

Although we often consider the management of business applications to be the pinnacle of systems management, there are many other scientific, academic, and government endeavors that require activities just as stringent as general business computing. Consequently, the five systems management disciplines are not aligned to any vertical market or work orientation.

Computing continues to be a general-purpose discipline where applied technology and operational practices migrate within the society of computer professionals who perform management operations. Practices that worked well in scientific areas have migrated into business over time, with both government and military agencies using unique database management, just to name two.

The point is that computer system management is an iterative process and continues to evolve through cross-pollination with different application bases, such as storage networking. Given that storage networking is a different storage model, however, existing management practices, though relevant, need to be molded to support a new infrastructure, which is what storage networking creates, a new application of storage.

In evaluating management of storage networking, we must apply the five disciplines. Figure 21-1 shows how these trans-technology practices are applied to the ever-increasing diversity within the data center further separating the systems into specialized practices. However, as Figure 21-1 also points out, storage networking draws heavily from data center practices and interfaces with each in order to be an effective infrastructure.

Planning computer systems to meet business continuity requirements has been one of the most tedious jobs within Information Technology. In the past 30 odd years of established data centers, business continuity responsibilities were ill-funded exercises in theoretical computer systems planning. However, as the world's socio-economic conditions began immersing computer systems within business operations, whether

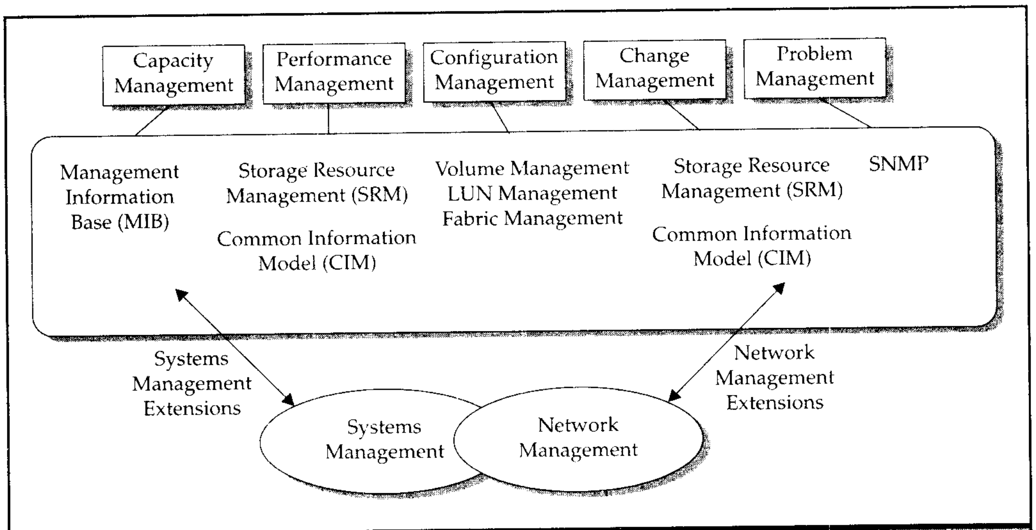
Capacity Management
Performance Management
Configuration Management
Change Management
Problem Management
Management Information Base (MIB)
Storage Resource Management (SRM)
Common Information Model (CIM)
Volume Management
LUN Management
Fabric Management
Storage Resource Management (SRM)
Common Information Model (CIM)
SNMP
Systems Management Extensions
Network Management Extensions
Systems Management
Network Management

Figure 21-1. *Fundamental systems management applied to storage networking*

for the sake of competitive positioning or national defense, the discipline of disaster recovery became a more imperative science.

Computer storage has always been the linchpin of any recovery scenario. From legacy systems to existing blade level embedded servers, all are linked to the vital circulation of data. Without it, the computer infrastructures they create would be worthless. Historically, the storage component was considered part and parcel to the holistic system that supported the evolving set of automation we have grown to depend on. Only recently has this link been recast into an infrastructure of its own. Storage networking in its inherent value has caused many to rethink their data center continuity plans, change their processes, and modify and enhance their solutions according to this new storage paradigm.

Defining the Environment (Putting Storage in Context)

The objective of business continuity is to facilitate uninterrupted business support despite the occurrence of problems. This depends on the problem, of course—it could be a small problem only effecting a single user application through an operative disk drive, or a complete system failure where the outage is caused by the entire SAN or NAS configuration being down. However, it could also be a site disaster affecting the entire data center. Certainly, the categorization of the problem is important, but the plans in place should be able to provide a road map, resources, and ability to hopefully recover from any incident.

Storage has traditionally played a key role in any recovery scenario. Without the data, the applications are useless. Consequently, most recovery scenarios center on making the data available as quickly as possible so business applications can continue the company's operations. Without stating the obvious, this requires the ability to replicate data throughout an infrastructure that enables everything from a micro recovery to a macro disaster site deployment.

Traditionally, with storage being tied directly to the server, this rendered storage problems and outages as associated problem instances—meaning they were viewed as a server outage and thus were combined with all the problems and outages associated with a server that's down or inoperative. As storage networking configurations create their own infrastructures, supporting more independent external components, they become active participants in the management disciplines. This redefines the data center environment and must be reflected in the continuity planning.

Figure 21-2 reflects the change from associated server/storage components to separate and discrete server, network, and storage components. This describes a data center that has divided its processing environment into three discrete infrastructures: servers, networks, and storage. Although many data centers have evolved into viewing and managing their environment this way, the separation of responsibility by operating environment (for example, UNIX, Windows, MVS, zOS, and Linux) must contend with the integration effects of heterogeneous storage networks.

Categorizing Interruptions in Service

The plans that are designed and implemented all depend on the magnitude of the problem incident. Therefore, it's important to distinguish and categorize the types of storage problems you are likely to encounter as well as larger problems that affect an entire infrastructure. This analysis provides a realistic view of the amount of insurance (redundancy, failover, and recovery methods) you need.

The interruptions in service that most commonly reflect outages can easily be lumped into two distinct categories: hardware failures and software problems. Hardware failures can be categorized through the inventory of storage networking devices and subcomponents, as well as the subsequent development of a matrix that reflects the causal relationship to each. Software problems can be organized the same way, although the causal relationships may assume complexities that are unnecessary to the task at hand. Because of this, it's best to be reasonable about what the system can and can't do when developing this list.

An enhancement to existing systems information will be to develop a matrix for SAN hardware components. Among necessary items such as component, name, vendor, serial number, and model, this type of list provides an excellent location to indicate a priority designation on the component in case of failure. Given our Plug-and-Play component-driven world, most problem resolutions fall into the category of a field

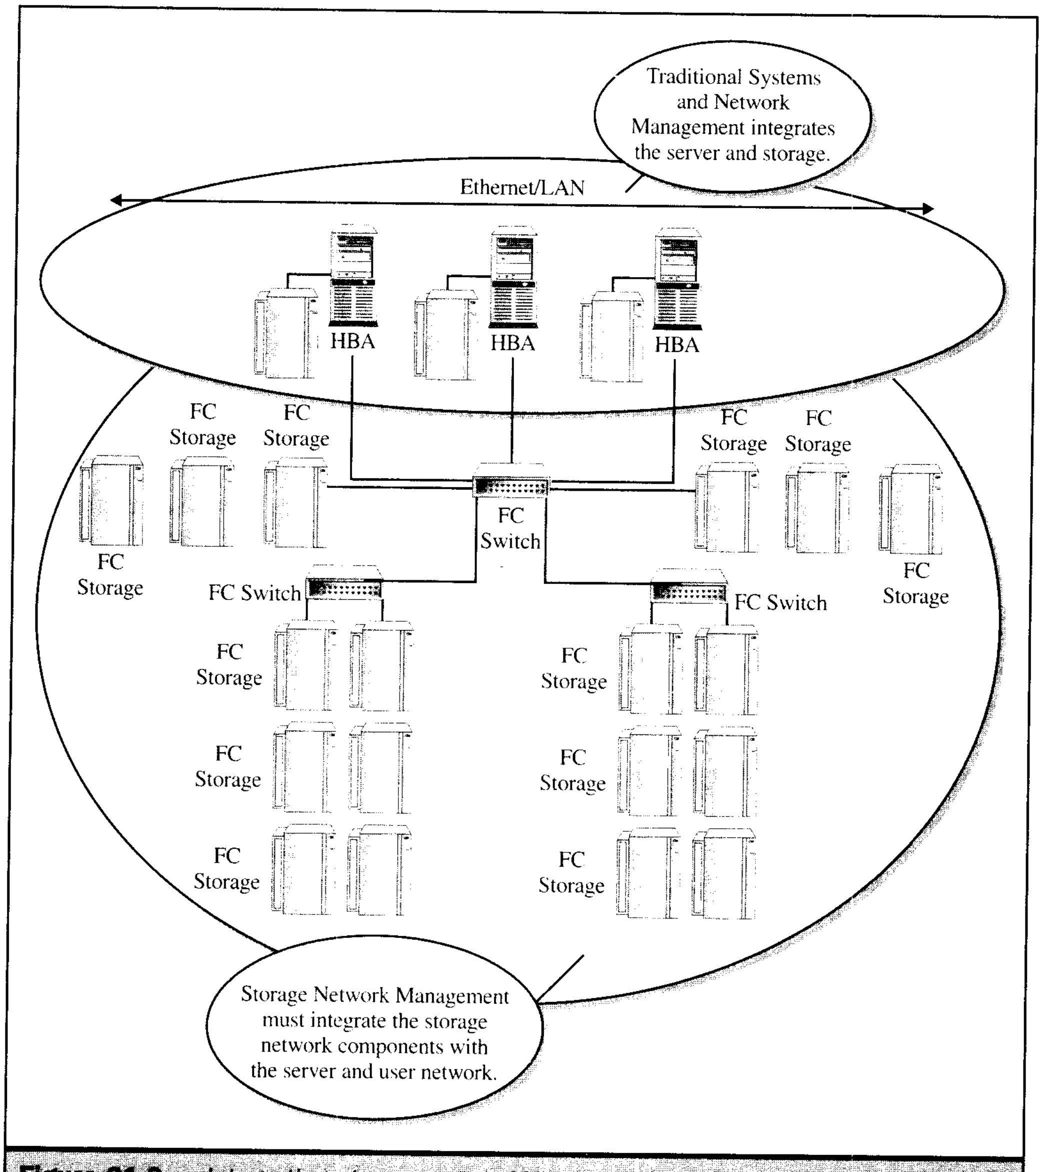
Traditional Systems and Network Management integrates the server and storage.
Ethernet/LAN
HBA
HBA
HBA
FC Storage
FC Storage
FC Storage
FC Storage
FC Storage
FC Switch
FC Switch
FC Switch
FC Storage
FC Storage
FC Storage
FC Storage
FC Storage
FC Storage
FC Storage
FC Storage
Storage Network Management must integrate the storage network components with the server and user network.

Figure 21-2. *Integrating storage networking as a major infrastructure*

replaceable unit (FRU), which provides an excellent beginning point to the categorization of the list. The value of having readily accessible and up-to-date SAN hardware information will enhance all management activities associated with the SAN hardware.

A corresponding matrix for SAN software configurations can be just as important. The relationship between switch operating system releases, node firmware, and server OS releases is necessary for problem identification and resolution. Keep in mind the SAN and NAS will place the overall solution in a multivendor scenario. The more information you have at hand, the better you can facilitate the problem resolution activities. Like the hardware matrix, the software list and categorization can articulate a level of priority and failure relationship to software elements as well as affected hardware components.

Work from the Top Down

Laying the groundwork for business continuity planning in this manner allows you to view the enterprise in a macro level. This enables the identification of a set of requirements to determine how a storage networking infrastructure fits into the data-center environment. This guides the level of planning that needs to occur at the micro level of storage continuity as you deal with storage recovery and, more specifically, the details surrounding storage networking.

If we refer to Figure 21-1, the ability to analyze the macro level is fairly simple. For example, go through the failure points from an externalization view first. This can be done, as shown in Figure 21-2, by performing a "what if" scenario on potential network outages. If the lease lines into the NAS servers go down, what applications go offline? Asking further questions qualifies the configuration for potential continuity insurance. For example, is the affected application or application service level compromised? If so, what is the recovery scenario?

In this manner, by working through the external scenarios, you can begin to develop a picture of potential fault points, their effect, and the basis for continuity insurance. In our example, with the leased lines being down, connecting any or all the remote offices would compromise their service level. Therefore, a level of recovery can be planned as insurance from this type of incident.

Participation

Planning and design of business continuity processes, resources, and environments require the participation of a diverse set of individuals. Although it's imperative to have the business users present, this may not always be the case. However, it should be the case that at the very least your colleagues in the data center participate, support, and even drive the plans and processes you develop. After all, the storage infrastructure is a key requirement for the continuation of automated business processes.

Without the participation and input from the business users you support, the plans you develop are considered "out-of-context." While that should not end your continuity planning, it puts the judgment of business application value and the assignment of recovery and redundant resources in making the storage infrastructure available totally within the IT organization. The out-of-context planning usually gets attention when the cost of business continuity is presented to the company.

The following includes additional detail and best practices when planning within your IT organization or within the user community:

- **Planning out of context** This provides a level of difficulty in determining the financial and operational impact placed on application outages, which may be accomplished through effective planning sessions with application designers and development organizations. It's best, as stated previously, to provide a robust continuity plan that encompasses all the production processing and related data, and then present this to company management with the associated cost. This will quickly lead to participation of the end users who ultimately have to bear the burden of these costs.
- **Planning within context** Planning continuity with the participation of the end users presents a much more realistic view of recovery and redundant requirements. Aside from a formalized project for business continuity, this should be integrated into existing capacity planning activities. Aside from formalized meetings, this offers a less intrusive method of gaining the information. This can be accomplished through activities such as worksheet-based surveys to end users that provide a macro view of the business impact and service level requirements.
- **Planning with our colleagues** At a minimum, the participation of your colleagues within the data center is imperative. As pointed out in dealing with out-of-context planning, the involvement of application design and development management should be involved to provide additional analysis on the value of the data and application service levels.

The Role of Storage Networking in Business Continuity

Storage networking, in changing the reality of storage deployment and access, shifts the traditional views of business continuity planning. Data becomes centralized, consolidated, and maintained within its own infrastructure (for instance, SAN and NAS configurations), and as a result, it has also become more distributed within remote areas of the company. Fundamentally, the change in technology requires different tools or at least a modification of traditional tools such as backup/recovery software, data replication, and archival tools and products. Regardless of these changes, many of the fundamental storage planning activities should continue to be applied to storage networking configurations.

The following is an example of the types of information that should be gathered in evaluating the business continuity planning requirements.

- **Data ownership** Data ownership continues to be more of a political question than a pragmatic one. However, finding the owner of the application data can provide the keys to the value, integrity, and necessity of it.

- **Application usage** Finding what programs access the data can provide the means to uncovering data that may be overlooked in continuity implementations. For example, there is always that one file that only gets accessed once a month by the inventory control application—*if* it's online. Consequently, not including this file in the regular backups, disaster recovery copies, or archival process results in an unhappy recovery experience.
- **Data utilization** This exercise provides a similar value to the application analysis. Not only does this uncover data that has not been accessed in months, quarters, or even years, it also provides a side benefit of increased storage utilization through the identification of unused data that can be archived to tape.
- **End users** Relating the data to an end-user department, division, or customer will also put another important piece of the continuity puzzle into perspective. Although this may seem clerical and administratively oriented, the proper placement of data within the complex structure of a SAN or the remote environment of a NAS are important to physical determinations of continuity configurations.
- **Current infrastructure** Understanding the current storage infrastructure, and applying the preceding bullets, provides a road map into what areas are most at risk. In terms of SAN and NAS, it will provide valuable insight into the placement of data from a NAS to a SAN and vice versa.

How Much Insurance Do You Need?

At some point, usually after you inventory what you have, estimate your requirements for continuity activities and resources. This exercise aids in assigning a value to the modifications in the current storage configuration installed, the additional redundant configurations you need, and other additional resources set for implementation in the event of trouble.

Level of Redundancy

How much of the production configuration needs to be replicated within what time period will be evident from your continuity analysis. In estimating the level of redundancy, it is necessary to understand the relationship to the minimum data capacity required, the number of users, and the level of access. For example, if you plan for full redundancy, then the data capacity, number of users, and access are simply duplicated in an adjacent configuration. However, this is rarely afforded given the cost, and therefore requires a level of adjustment that takes into account the type of outage, and the level of business operation required for a particular time period.

For example, for unscheduled outages, the level of operation must be maintained for deposit transactions, with the related data available. However, the other transactions may be down for a period of no more than 15 minutes for a single processing cycle.

Therefore, configurations have to reflect a level of redundancy to maintain the deposit suspense transactional data while the other transaction data, customer service, and portfolio services should have fault tolerant systems with a mean time to repair of 15 minutes per 12 hours.

Level of Requirements—How Much It Will Cost

In calculating these requirements, you must not only account for new configurations (there to meet necessary service levels), but also those modifications—reflected in existing production configurations—that accommodate redundant systems. In terms of NAS, this may be the addition of a redundant NAS device that mirrors the data. Within the more complex SAN configurations, requirements include: upgrades to switches to add expansion ports, E_Ports, software to enable interswitch linking, trunking, and redundant switch configuration and related devices.

Level of Associated Requirement—How Much It Will Really Cost

One of the most challenging activities within this exercise is the inclusion of additional hardware and software components you didn't think about initially in contemplating storage configurations. These expenses include additional software licenses, upgrades to existing software, and the cost of new devices such as adapters and redundant servers.

Storage Design and Implementation of the Business Continuity Plan

As you develop a new plan for the business continuity of storage networks, or modify an existing plan, there are key areas to consider that will enable the network to continue operating in a variety of circumstances. Among these are the integration of business continuity requirements into capacity planning disciplines, the effective tracking of storage system and device outages, and a realistic evaluation of sustained outage assessment.

On the technical side, the ability to effectively monitor performance and capacity of existing configurations, though challenging, is imperative (see Chapters 22 and 23). This includes an inventory and evaluation of the existing capabilities of storage functions (for example, RAID level support, environmental limitations, and remote settings).

SAN

The main design considerations for the SAN are the configuration of core switching and estimates regarding new redundant switches for recovery processing. Within this category are the number of expansion ports, the availability of dynamic ports and paths, and port utilization estimates. Secondary to the core switch is the storage strategy to facilitate replication of recovery data, the time and state requirements of replicated data, and the synchronous (or asynchronous) nature of this process. Thirdly, is the ability of the reconfigured infrastructure to handle the adjusted I/O workload.

It should be noted that many additional items are required to develop, design, test, and implement a full business continuity plan for SAN. These include server sizing, recovery site preparation, application portability, and required application software and subsystems, such as compatible file systems and database systems. Although outside the scope of this book, these items should be integrated into a well-designed plan and not overlooked.

NAS

NAS will be the most deceptive in preparing for business continuity planning. Given its ease of installation and maintenance, it's possible to overlook its potential problems. Design considerations should center on the requirements for data redundancy, and result in configuration to mirror production data either synchronously or asynchronously. This requires a NAS-to-NAS device configuration and should be driven by the requirements for complementary or compatible devices.

In addition, it will be necessary to provide redundant capacities and daily backup protections for respective devices. Over and above these are fundamental considerations to ensure that capacities are allocated so that as a failover scenario is implemented, the realities of the continuity plan workload can be handled.

Remote vs. Local

Redundant configurations that are local must adhere to cabling limitations. Remote configurations must maintain a strict level of monitoring to ensure the data maintains the prescribed level of consistency. Local redundancy costs can be much less than remote—however, full failover clustering systems can move costs well into the range of remote disaster recovery services.

Backup vs. Disaster

It is important to understand whether the configuration will support a local system or unit outage, or have to be part of an entire disaster recovery scenario. This simple characterization of the business continuity solution will result in a configuration direction focused on the appropriate support.

Expertise and Configuration Data

Trying to not overlook necessary components in assembling an alternative storage configuration to support a business continuity solution is a constant worry. However, the successful operation of disaster recovery storage configuration may never see production processing and hopefully never be used in a disaster scenario. Even so, these are configurations that certainly should be tested. A successfully tested configuration should provide the required information of configuration settings, files, and addressing schemes including storage arrays, switch ports, LUN assignments, and zonings used.

Many components of the backup storage infrastructure may have to be re-configured back into a production environment or used in an alternative configuration. If this is the case, or might be the case, be sure you document the configuration information for later access. It's best to rely on procedures and current information than ad hoc expertise to reinvent the backup configuration during each disaster recovery test.