Chapter 23

Maintaining Serviceability

The failure of a switch port, the inoperative disk drive, and the new fabric software all require service to the SAN. Upgrading the NAS storage array, the need to add Ethernet ports, and the installation of the new NAS device all require a change to the storage network configuration. Each of us is familiar with the call at 2:00 A.M. when the production application is down and problems point to the SAN or a suspected NAS device. These are definite problems and require an effective problem identification, tracking, and resolution system.

Servicing is defined by the activities necessary to maintain the hardware and software configurations of computer systems. Storage networks like other infrastructures require servicing. SANs, still being a relatively new technology, are more intensive than most for many reasons, not to mention the diversity of components that make up the configuration. NAS configurations, on the other hand, being somewhat more mature, require very little servicing, although this depends on their application and usage volatility. Regardless of either application of storage networking, each storage infrastructure benefits from a servicing system that provides significant configuration reliability and, subsequently, availability.

Servicing is often driven by problems that occur within the storage network configuration as well as by external problems that affect the performance of the storage network. Secondly, servicing is required for any changes to the configuration meant to fix problems and to enhance or upgrade configuration components. Considering another oft-used work axiom—*if it ain't broke don't fix it*—the necessity of using a variety of sometimes-compatible hardware and software components brings another dimension to this colloquialism.

From a systems management discipline, the next phase of activities are change management, problem management, and, as its cornerstone as indicated in the following illustration, configuration management. Each plays a pivotal role in the effective servicing of the storage network. Traditional storage management and elements of network management make up many of the elements that are necessary to service the storage network. This requires an integration of traditional configuration management and change management activities. Not to mention the effective development and integration of problem management into the data-center operations and help desk activities.

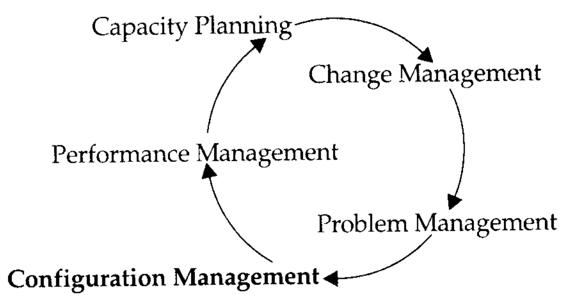

This chapter discusses the application of these management disciplines to the storage network environment. This includes the complexities of hardware/software compatibility issues that affect SAN configurations and the added value of keeping NAS release levels in sync with each other and the network software. Within the data center, one of the many technical issues to deal with is the integration into other system management disciplines that are already in place. Of these, problem management is one of the most challenging.

The activities to track changes within the storage network are beneficial in order to keep the devices compatible, correlate problem identification, and schedule maintenance activities. However, the benefits of effective change management within the storage network also provide value to external elements within the data-center environment. Integrating yet another level of change into the data center is a challenge. This chapter provides some insight into these potentially volatile situations and offers some guidelines to get things started.

Tracking the Configurations

Computer systems require the installation of multiple components that, when working together, form some type of system for producing effective results. Generally, those results are productive for the IT organization responsible for the system. The ability to manage the diversity of hardware and software components is referred to as configuration management. Configuration management establishes the inventory of a system, naming the components, their release and maintenance levels, and connectivity. Although configuration management is suited for both hardware and software, it is usually used in a segregated manner to track the hardware configuration and the software release levels, respectively.

Although the premise is that by keeping track of the components within the hardware or software configuration one can understand the economics in terms of tracking capital equipment, the alternative but oft-unused value of this discipline is the tracking of the inventory to understand what changes can be made, under what circumstances, in order to either correct a system deficiency or enhance the performance of the system to meet established service levels. Given that both hardware and software are likely to change over the effective life of a system, configuration management serves as both a driver to changes and a tool to identify problems.

Storage networking, being that it's a system, and a relatively new system at that, can benefit greatly from configuration management. Establishing an effective servicing practice for storage networking systems, and keeping track of the diversity of components that make up SAN or NAS infrastructures, can be extremely valuable.

What's different about SANs in configuration management is the number of separate components that make up the configuration. In addition to its diversity, the shared characteristics of both networking and storage technologies gives SANs a hybrid view to many people uninitiated in storage networking. At the same time, the SAN infrastructure is driven by the external effects of other compatible and collaborative systems, such as servers, networks, and software.

NAS poses a more chameleon-like problem. NAS configurations can be single or multiple in number, and while looking like storage, they are really a specialized server. This provides a configuration management discipline with some interesting problems in determining where to categorize this component and what elements make up the components. If it's a storage device, then it has no OS. However, NAS most definitely has an OS that needs upgrades, changes, and has the potential for problems. However, tracking as a storage device overlooks these critical elements. In addition, NAS configurations are driven externally by the network and remote servers or clients that contact it through remote communications.

The point is that storage networks, although computer systems in and of themselves, are different animals when it comes to categorizing them through traditional configuration management practices. Some thought must be given to these configurations as they are integrated into either existing configuration management tools or are used in establishing new practices that integrate both an inventory tracking for capital expense purposes and a foundation for tracking problems and change.

Physical Configuration Management for SANs

Configurations are likely viewed from a hardware and software perspective. Physical views of the SAN should provide the inventory of all equipment installed within the configuration. This requires all the typical information necessary for asset and capital inventory tracking, such as serial numbers, location, and ownership. In addition, there should be the same current information on firmware or micro-code releases and install dates. Many components in the SAN area come with some type of firmware, micro-code, and micro-kernel software that is release- and hardware-specific. It's important that this information is collected and stored, at a minimum, on a file, although a database is better for later access.

Thus, a decision will have to be made here in terms of categorization of software. Given that firmware and micro-code are hardware-specific, the micro-kernel that functions as the fabric OS must have a place here or in the software configuration tracking system. Consequently, the choice depends on the size of the configuration and what type of distinction there is between hardware and software support for servicing. In a large installation, this could be critical when it comes to avoiding confusion in tracking microkernel release levels in two places, or dodging confusion with other firmware and driver software that is specific to HBA and storage array controllers.

Figure 23-1 indicates the types of categories that may be considered in establishing the configuration management repository for SANs. Of course, this information can and should drive the digital rendering of the configuration for general reference and access.

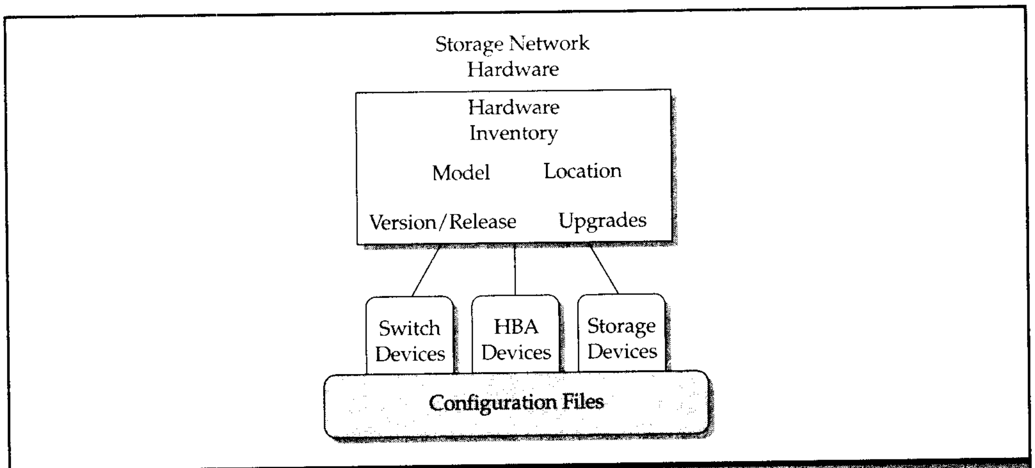

Figure 23-1. *SAN configuration management categories*

Logical/Software Configuration Management for SANs

The logical configuration refers to the SAN software, but also should include the logical relationships that exist within the SAN configuration. These are the fabric zones, the LUN assignments at the server/HBA node, and the storage controller array nodes. In addition, there are additional elements within the fabric operating system that relate to the port IDs, fabric logins, worldwide naming services, and the management information base dictionary and layouts.

Configuration management for the logical and software functions is critical. Due to the diversity and complexity of the SAN software components, much time can be saved and downtime avoided through the effective documentation and monitoring of the fabric micro-kernel release and maintenance level, the HBA driver code release and maintenance level, and the storage array RAID driver and firmware release level. There should be an established configuration management for the attached servers. If not, the server OS release and maintenance levels will be mandatory.

Essentially, the software components of the SAN form a matrix of interdependent functions. A change in any has the potential to cascade problems throughout the SAN configuration. You can rely on interoperability labs of vendors, industry groups, and testing labs, to gather macro level compatibility information. However, the number of permutations far exceeds what these well-intentioned groups can provide. If that means providing a test bed for your applications within the configurations specific to your site, then it's well worth considering. In fact, it's no different than other test beds that you may require for application and system testing and should be integrated into testing facilities.

Figure 23-2 illustrates the configuration management categories and example specifications a data center may consider for SAN physical and logical components. The interrelationships and dependencies between components can quickly become overwhelming, which provides additional justification for building a repository of physical and logical components. However, one of the most difficult events to track are the dependencies from hardware to software components. Consequently, the ability to begin an integrated repository may save critical time as configuration changes can be tracked against the related hardware and software device dependencies.

Physical Configuration Management for NAS

Managing the physical configurations for NAS might seem like a simpler process than SAN, but it's not. The physical connections within the NAS configurations may be simple if they are installed in close proximity and operate within the confines of the data center. Even then, it can become confusing and a challenge because of the diversity of networks and network components. The network becomes the critical factor in establishing the configuration, and consequently, the NAS configurations pose a relatively difficult problem in regards to configuration management—that is, how to articulate and track individual components of the device and the external characteristics of the network hardware within the context of a system (in other words, the infrastructure).

Figure 23-3 illustrates the complexity of the external relationships between the network, the NAS device, and the attached servers. In a sense, it can be viewed as a loosely connected storage network, with dependencies associated with network topologies and protocols, server nodes, and file communications protocols. Unless the hardware and firmware release levels are inventoried and tracked in conjunction with the network, the NAS systems become unassociated storage servers unbound to the confines of the networks in which they operate. Bottom line, if there is a problem

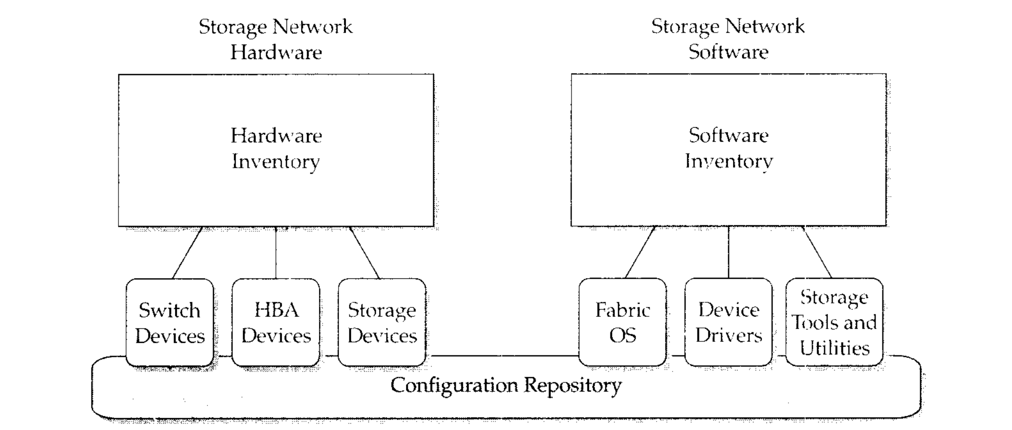

Figure 23-2. *Integrating SAN physical and logical categories into a single repository*

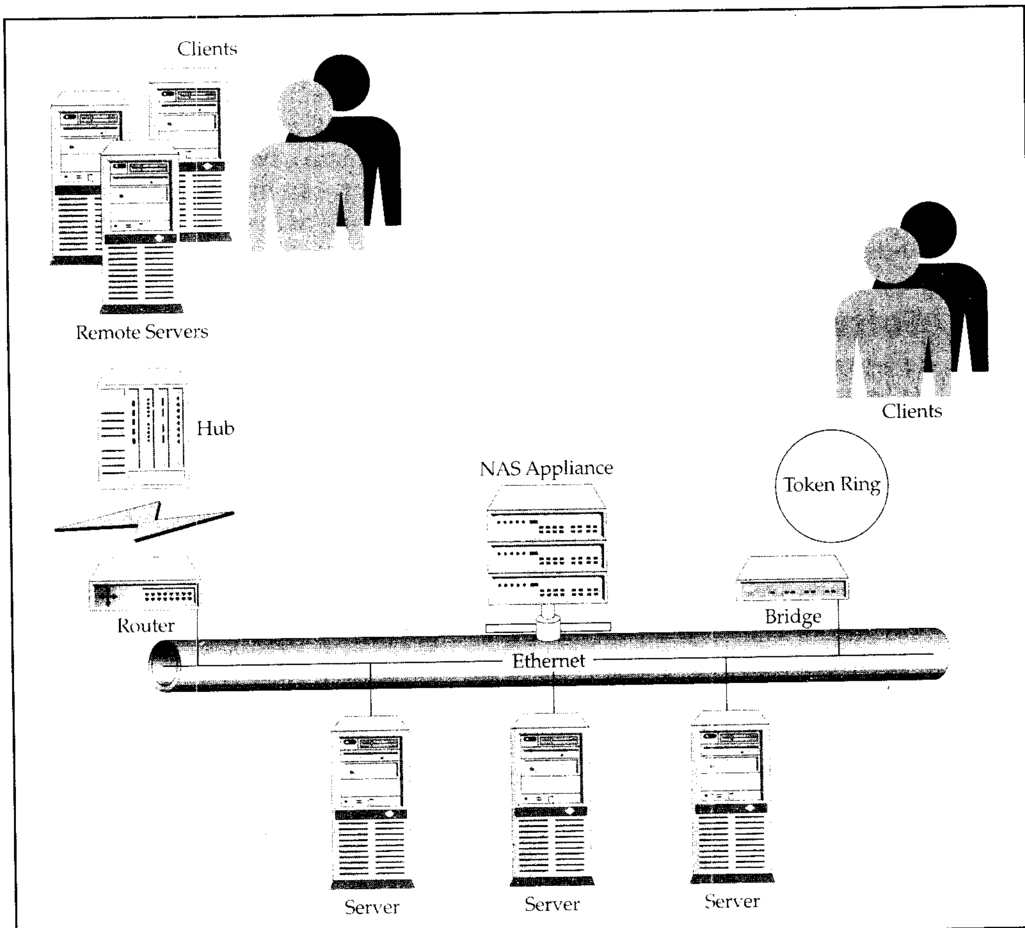

Figure 23-3. *The complexity of NAS external relationships*

within the NAS device that's being perpetuated through the operations of a network device, the problem will be hard to diagnose and will likely repeat regardless of any changes within the offending NAS device.

Consequently, it's paramount that some type of relationship be established as the configuration management is instituted for the NAS configurations. This may require the NAS device being subsumed into the network configuration management disciplines, if they are established; however, this may not take into account the other relationship of the server component that is actually generating the I/O request. Given this third variable, the NAS configuration management must rely on two external factors for effective configuration management: the network and the server. Whether these are accounted for in the network disciplines or the system disciplines or a new storage network discipline, the benefits to problem investigation and upgrade analysis, as well as overall availability, are large.

Figure 23-4 indicates the types of categories that may be considered in establishing the physical configuration management matrix for NAS. Note that due to the complexities of the external relationships of the network devices, links into a network configuration repository enhance the value of the NAS physical configuration information. Of course, this information can, and should, drive the digital rendering of the configuration for general reference and access.

Logical/Software Configuration Management for NAS

Software for NAS becomes both simple and problematic. Simple in terms of the standardization of the NAS micro-kernel OS used, given that the configuration is a homogeneous vendor solution. Problematic in terms of how the NAS OS reacts to a network environment of heterogeneous components (for example, hubs, routers, and switches). Coupled with this is the reaction to a heterogeneous environment of server OSs that may cover both UNIX and Windows variants.

Given that the NAS micro-kernel forms a barrier to direct communications from remote clients to the storage arrays through file access, the storage arrays associated with the NAS solution are bundled and protected from extraneous and external access. The NAS storage upgrade path may be non-existent on entry-level devices and limited

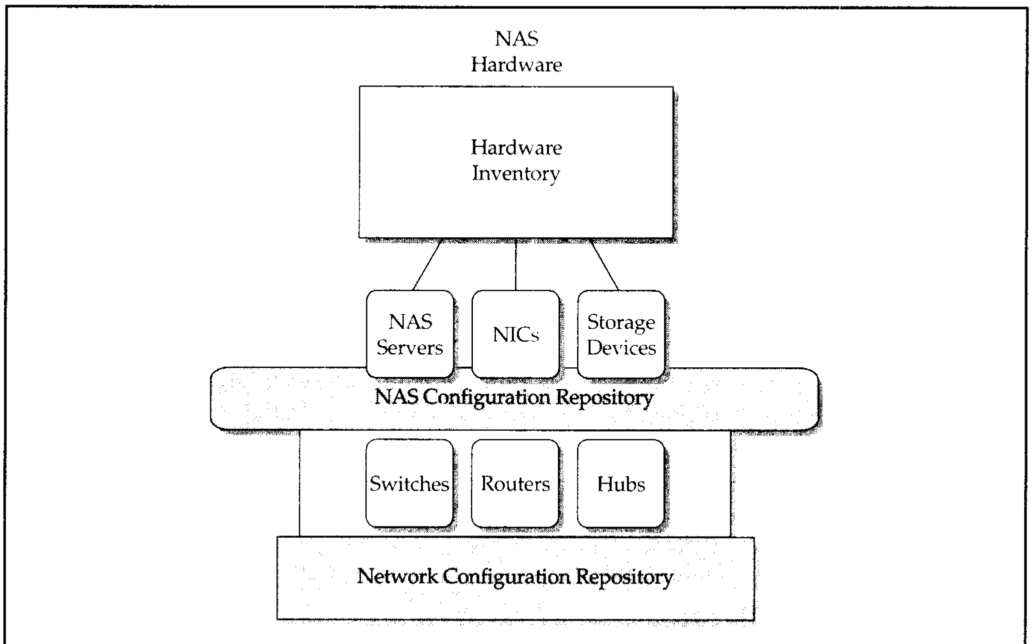

Figure 23-4. *NAS configuration management categories*

to specific vendor upgrades with larger enterprise-level devices. The RAID controller firmware will thus be important to problem investigation associated with vendor servicing.

Configuration management within the NAS infrastructure must encompass the NAS micro-kernel release and maintenance level, the storage RAID controller firmware release and maintenance level, and the network driver release and maintenance levels. These fundamental parts can be associated with network software, network components, and server OS release and maintenance levels to form an effective NAS configuration relationship matrix.

The logical constructs of NAS configuration can be valuable in micro views of the data storage strategies used in the NAS devices. In other words, these are the logical views of RAID deployment within the arrays with configurations like local mirroring and remote mirroring illustrating the synchronization methods used. Integrated with the logical storage information, the logical views of file system extensions supported by the NAS device should be documented.

Figure 23-5 indicates the types of categories that may be considered in establishing the software and logical configuration management matrix for NAS. The complexities of integrating both hardware and software should be taken in context with other configuration management files, repositories, and tracking mechanisms. Figure 23-5 suggests the linking of not only the network inventory but the systems inventory of application and database servers and related information. Of course, this information can, and should, drive the digital rendering of the configuration for general reference and access.

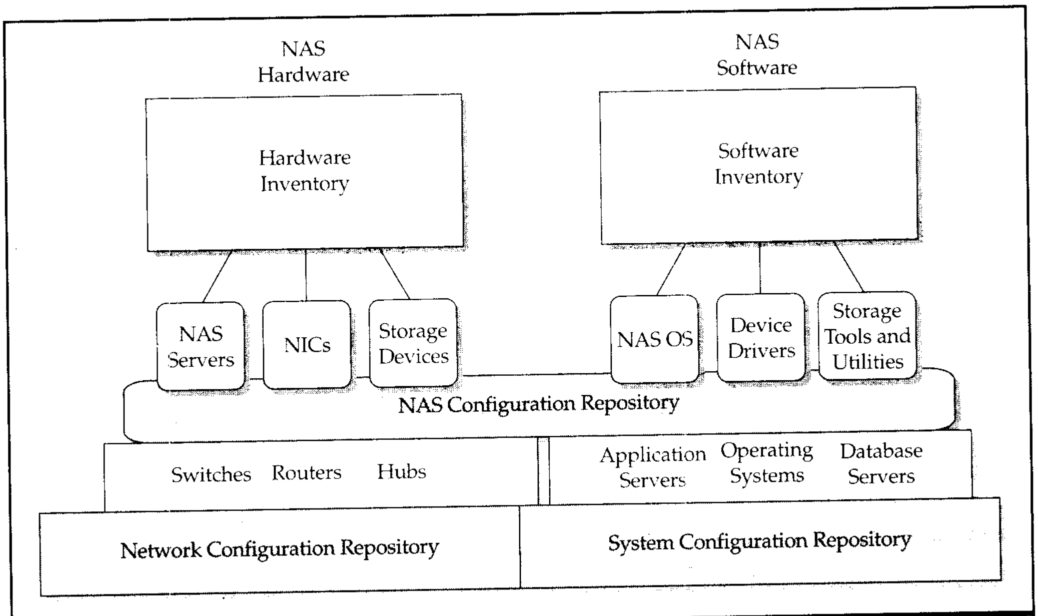

Figure 23-5. *NAS software and logical views with example categories*

Investigating the Changes

What is change management and integrating SAN/NAS changes into an existing data center CM system?

When was the last time someone made a change to the configuration you were supporting? Don't know? Well, probably about five minutes ago, if you're operating in an unregulated environment. The chief creator problem is change. Something changed, and finding out what changed constitutes the bulk of problem analysis. Finding who changed what becomes more of a people management issue than technical. However, the discipline of change management within computer systems came about for just these reasons. Who changed what, when, and where is the mantra of change management.

If tracking and prioritizing changes is already an established discipline within the configurations you support, then you'll understand the value of integrating an unwieldy infrastructure like storage networking into the change management activities. The central value of managing change is that in the case of young technology solutions like SANs, almost 90 percent of the problems can be eliminated just by understanding the configurations and implementing reasonable changes.

The chameleon-like nature of NAS and the ease and attraction of Plug-and-Play provide a tremendous force in circumventing the change discipline. Well, you say, it's just some network storage and I can put it on the network in less than five minutes. Yes, you can, and you can also bring the network to its knees, lose all the data on the NAS, and corrupt the data on the redirected server. Wow, that's a lot of potential problems for Plug-and-Play. Regardless of whether we're accustomed to this type of service, it doesn't have to happen.

Forget the formalization of change committees and change windows for the SOHO environments. All you need is some common sense to discipline your changes. Developing a change management discipline can work just as well for small environments as for large environments. No one appreciates the use of the defragmentation process that runs during the peak of the server's day, especially when it crashes and all the sales transactions have to be reentered because the mirror on the NAS device didn't work. Consequently, even if you're not familiar with change management, starting with storage networks can be extremely beneficial, all in the name of availability.

Driving Change Through Configuration Management

An ongoing program of regular monitoring and analysis of the storage network configuration ultimately reveals changes that must be made to enhance performance or provide sufficient reliability to meet service levels. Whether this comes through changes in software or hardware activities, the ability to schedule, track, and implement these changes effectively, and in a non-disruptive manner, requires the disciplines of change management.

In many cases, however, there will be changes to the configuration that require a timely response. Consequently, there must be a prioritization within the change management system to handle these issues. In terms of SAN serviceability, the inoperative switch that impacts the service levels of a production application system must be analyzed given its relationships to the entire configuration and its effects on other components and systems being supported through the SAN configuration. This may require a quick response, but without sufficient and accurate knowledge of the existing configuration, the ability to identify and resolve the problem becomes difficult. Either the change is made without knowledge of its impact, or a lengthy analysis to document the configuration may elongate the outage or degrade the service.

Consequently, configuration management drives changes by documenting storage network "hot-spots" and quickly referencing the relationships within the configuration. In addition, an effective change management program can leverage the information within configuration management system diagrams to facilitate the problem resolution process in conjunction with change prioritization. Helping to facilitate the establishment of change prioritization is an added benefit of configuration management that drives change management in a positive direction.

How Problem Management/Resolution Becomes the Ying and Yang of Storage Networks

You can't fix a problem unless you change the configuration. You don't have problems unless there is a change to the system. This is the "yin and yang" of change management. Although we have yet to discuss problem management, we have described the effect of changes within the configuration being driven by problems that occur. Because storage networks are implemented using diverse components, they are exposed to a disparate set of external stimulus. They become the critical elements to production systems, however, and as such, they are the prime focus when diagnosing severity level one problems. In other words, storage is a key element to evaluate during system outages.

Consequently, managing change can drive the balance of change within the overall data center activities. After all, it's easy to blame everything on the new kid (say, storage network configuration), if they can't defend themselves. The need to enroll storage network configurations into change management activities allows many things, including the visibility of changes within the configuration, visibility to external changes that affect the configuration, and the ability to maintain a credible defense in problem identification activities.

The Effects of Capacity Upgrades on Change Management

One of the major factors that disrupt storage configurations are upgrades. Whether there're extending capacities or enhancing performance, these changes should affect existing service levels in a positive way; however, in many cases, capacity upgrades can provide unforeseen results and performance, generally in the wrong direction. These situations evolve from unscheduled or unmanaged change plans that do not account for adequate time to complete or consider the effects to related systems.

This is evident in the installation of additional storage capacity using NAS devices. Without sufficient consideration for the related network capacity and existing traffic, NAS devices can solve the storage capacity problem but create a network problem as a result. The other issue with NAS upgrades is the related file protocol services that need to be connected or reconnected after an installation. Installing or upgrading NAS device capacity may be quick in terms of configuring and identifying the network; reestablishing the drive mappings and device mounts for the new NAS devices are sometimes an overlooked detail.

SAN configurations provide a more complex activity for upgrades and capacity upgrades. Due largely to the diversity of components and their interrelationships, SAN upgrades and installations will require a significantly longer planning and analysis cycle than other storage configurations. For example, upgrading storage capacity has a ricochet effect within the configuration. Adding a significant amount of capacity is driven by increased access. The increased access may require additional node servers to be attached, and ultimately will place more traffic within the switch configurations.

If the extenuating situations as mentioned in the preceding example are not addressed, the SAN configuration can become compromised in the name of increasing service. Addressing these situations requires a discipline that is supported through a comprehensive change management system. The articulation of changes within the storage infrastructure drives the discipline to analyze the following: what's being changed, why, the desired result, and the plan for reestablishing service. Managing changes highlights the division of labor between software and hardware changes. It also highlights the division of responsibility in large configurations where the network part and the storage part are separated and supported by different groups, each having their own hardware and software expertise. Integrating change management into the storage configuration management process provides a forum to highlight all the required activities for servicing the SAN or NAS configurations.

Tracking Changes to Physical and Logical Configurations

Tracking of changes within a storage network is similar to the tracking of changes within the data center: a separation of hardware and software changes. Another view is to divide them into physical and logical categories. This mediates the traditional separation of hardware versus software, thereby integrating the component responsibility into combined or shared responsibilities. For example, the typical enhancement of storage capacity within a storage network is the installation of a new and larger storage array. This could be performed in an off-shift time period, or on a weekend, and no one would be the wiser.

However, it could also be scheduled within a change management system where the information regarding the upgrade, proposed activity, and scheduling are communicated to other areas that may be effected. This allows others to either consider the effects of the change and plan accordingly, postpone the change until issues are addressed, or, at a minimum, be aware that the possibility of problems could exist after the change date.

Because both NAS and SAN configurations are closely associated with hardware and software functions (see Parts III and IV for more details on the major components of these storage systems), the tracking of physical changes provides a method that encompasses both the physical aspect of the component as well as the bundled software configuration parameters that need to be addressed. For example, the recabling of a SAN switch to enable the installation of a new switch requires both physical changes and subsequent fabric changes. Viewing and administrating changes in this way forces the change process to encompass both hardware and software skills while ensuring the changes are analyzed in an integrated fashion.

The tracking of logical changes performs the same function within the change process, except that logical configurations are generally driven by changes to software components within the storage network. These may be enhancements to various features and functions, or maintenance of existing software elements such as the fabric OS, NAS micro-kernel, and storage and adapter drivers and firmware.

An example is the upgrade to the NAS micro-kernel software. This change, viewed from a logical perspective, should require an analysis of configuration management information regarding which relevant NAS devices are affected by an upgrade. The NAS target configurations would then be evaluated for subsequent firmware changes within the storage array controllers that may be required to support the new NAS kernel software. In addition, the micro-kernel change may have different levels of IP network support and will require the evaluation of the potential network components. Consequently, the change to the logical configuration brings in all of the NAS components, as well as the external elements that may be affected. This also forces the necessary storage and networking hardware skills to participate in the scheduling and implementation of the change.

Guidelines for Establishing Change Management in Storage Networking

The following guidelines can be used to provide the beginning of a change management program for storage networking. Even though a formal change management function may exist in your data center, the establishment and integration of storage networking as a valid change management entity is key to participation at a data-center level.

- *Establish a link between storage configuration and change management.* This forces the evaluation of each change to be analyzed for the entire configuration, as well as other systems in the data center.
- *Develop a relationship matrix between the key components among the storage network configurations.* For example, the relationship between the FC switch and the storage nodes is a fundamental and interdependent one. If the switch becomes inoperative, the storage is inaccessible and potentially magnified by several storage arrays supported by the switch. Using this example, it also provides a roadmap of Inter Switch Links (ISL) dependencies and other fabric nodes that are participating in the fabric.

- *Establish an internal (within the storage network) change prioritization scheme and an external (outside the storage network) change effects scheme based on a storage network availability matrix.* The storage change prioritization scheme will define the levels of change based upon their potential effects on the storage configuration. The external change effects scheme will define the levels of how the storage configuration is affected by outside configurations of servers, networks, and other storage configurations.
- *Establish a link between storage change management and problem management.* Additional discussion on problem management follows; however, it's important to mention the critical link between problems and change given their synergistic relationship. For example, if a storage configuration denotes a FC switch that is attached to five storage arrays and another two switches where information drives the relationship matrix that depicts its interdependencies. These component interdependencies drive the change prioritization scheme to reflect this component as a critical availability element within the configuration. The external changes affect change prioritization schemes that, in turn, reflect the risk to external systems supported by these components. Consequently, this will drive a level of change prioritization if and when a problem arises within this component.

Whose Problem Is It Anyway?

Storage problems within the data center can ricochet through production applications faster than any other component. Obviously, this is a reflection of their critical position within the diversity of components that support production processing. When problems occur, there generally is some type of trouble ticket system that logs problems for assignment and potential resolution. These systems provide an invaluable service when it comes to documenting and tracking problems, but does little to facilitate the entire aspect of problem management. Consequently, applying a trouble ticket system or help desk software to storage problems only addresses a small portion of problem management.

Problems, especially within the storage infrastructure, are resolved in a number of ways, depending on the data-center facility, administrator preferences, and management, or even the establishment of service levels. Regardless of data center–specific attributes and management styles, there are fundamental activities that reflect problem management disciplines as well as critical activities that must be performed to respond to storage problems effectively. The problem management disciplines are: problem prioritization, problem identification, and problem resolution. Overlaid with problem documentation and tracking, they form the fundamental parts of problem management. How these are implemented and integrated with storage networks is critical to the success and reliability of the storage infrastructure.

The Problem Prioritization Debate

There is no debate more dynamic within the data center than that surrounding problem prioritization. Someone must decide what constitutes a severity level one problem, as opposed to severity level two, three, and so on. This discussion becomes even more vigorous when dealing with storage, since storage contains the data, and ownership is an attribute of that data. If the data becomes unavailable, then the application that uses the data is unavailable and the level of impact is driven by the effect this situation has on the activities of the company. Consequently, end users drive the establishment of severity definitions with the data availability, vis-à-vis, storage availability. However, most are ill-equipped to handle or understand the level of complexity within the storage infrastructure to effectively provide this.

These discussions will become either a very complex issue or a very simple equation given the level of data ownership taken by the end user. If end users take data ownership seriously, then discussions regarding the definitions of data, usage, and management will become involved. However, if end users simply see the data center as the service supplying the data, the discussion can be reduced to a simple level of "on" or "off." In other words, either the data is available, or it isn't.

Regardless of these two perspectives, and their highly charged political engagement with end users, the data center must view problem prioritization in a more utilitarian fashion. Simply put, the complexities of providing data delivery and availability infrastructures must be managed within a matrix of interconnecting systems that are beyond the level of knowledge or interest of end users. Moving end users toward a utilitarian approach allows the data center to deal effectively with storage ownership and management of the infrastructure. Defining problems from an end-user perspective within this increasingly complex set of infrastructures allows data centers and underlying areas to define the prioritization required for the data utility. Key to establishing prioritization mechanics is storage networking.

The relationship between data availability and storage networking problem prioritization increases the complexity of this activity. Given the characteristics of data delivery and the evolution of decoupled storage infrastructures such as storage networks, the chances that problems will arise, given the level of interrelationships within the storage area, increases dramatically. This results in environments with increasingly complex storage problems while trying to maintain a simple utility data delivery system.

Problem Identification

Investigating and identifying problems within storage networking configurations is an ongoing activity—for three reasons. One, these technologies are still new and continue to be temperamental. Two, storage networks consist of multiple components that must interact both internally within themselves and externally with other components and systems within the data center. And three, the industry has not yet produced the level of tools necessary for effective operational problem evaluation and identification.

However, there are some key items to consider when looking at problem identification in these areas. Some activities don't require tools so much as basic problem-solving practices. Although this is beyond the scope of this book, the necessary activities should be mentioned in context with storage networks.

- **Accurate and Current Information** Configuration management for installed storage networks should be accurate and readily available.
- **Diagnostic Services** A minimum level of standard operational procedures should document the entry-level activities for problem investigation and resolution. This should include the dispatching of vendor service personnel in order to facilitate problem resolution.
- **Service Information** Accurate and current information should be available to contact vendors and other external servicing personnel.
- **Tool Access** Access to consoles and other logical tools should be made available to support and operations personnel. This includes current physical location, effective tool usage, diagnostic processes, and authorized access through hardware consoles and network-accessible tools.

In SAN environments, an effective guideline is to investigate from the outside in. Essentially, this requires the evaluation of problems that affect all devices versus problems that affect only specific devices within the storage network.

For problems that affect all devices, it's necessary to view components that are central to the network. Specific to SANs is access to the availability of the fabric operating services. The directly connected console is the most effective tool if network access is not available, while fabric configuration routines provide status information regarding port operations and assignments. There may be additional in-band management routines as well, which can be used for extra information.

For problems that affect only some of the devices, it becomes necessary to first view components that are common to the failing devices. If the total SAN configuration is not completely down, then out-of-band management routines can be accessed to provide a port analysis, port connections, and performance checks. In addition, the ability to telnet or peek into the attached node devices and node servers verifies the basic operations of the switch fabric.

Specific to the node devices are the mostly out-of-band tools which evaluate and verify the operation of storage devices, server HBA ports, and server operating systems. However, access through the fabric can also verify connectivity, logins, and assignments. Key to this area are the continued network drops from the fabric by node devices, a practice which can be traced more effectively through fabric access using configuration and diagnostic tools.

In NAS environments, the same approach can be used to identify failing components; however, NAS configurations are more accessible to typical performance tools. Therefore, the most effective approach is to view problems from the outside in by focusing on access points common to all NAS devices—for example, the network. This is followed by the investigation of individual NAS devices and related components that are failing.

The centralized approach can be used with existing network tools to monitor and respond to NAS network drops where access within a specific area may be the key to the problem (for instance, a network problem versus a NAS problem). This demonstrates how NAS devices require less overhead to identify problems because of their bundled self-contained status and their integration into existing networks as IP-addressable devices.

If problem identification focuses on an individual device, then the failing components may be simple to find given the bundled nature of the NAS solution and the available tools within the NAS microkernel operating system. All NAS vendors provide a layer of diagnostic tools that allow investigation of internal operations of the device. This becomes a matter of a failing server or storage components given the simple nature of the NAS device.

NAS devices are less expensive and easier to integrate into existing problem identification activities. Given their simple configuration and mode of operation as an IP addressable network component, they can be identified quickly through network tools. In addition, internal diagnostic vendor tools provide performance statistics and notification of failures.

Problem Tracking

Once a problem has been reported, the ability to document its particulars is critical to the identification process. Most data centers will have an existing problem documentation system associated with their help desks. Integration of these systems is important for setting the foundation for problem tracking within the storage network. This doesn't automatically mean that all problems will be documented through the help desk, however. Many storage network problems will come from internal sources where they're tracked and reported by data-center administrators, data base administrators, and network administrators, to name a few.

Even though these positions may report the bulk of problems, it does not guarantee those problems will be reported through the same trouble ticket system used by help desk and operations personnel. Thus, the foundation is laid for problem duplication and multiple reporting. This demonstrates just one of the challenges regarding implementation of the storage network into the problem tracking system. The same can be said for other key areas of the problem reporting process when it comes to critical and accurate reporting of problems.

Storage network problem documentation is made more effective by following a few general guidelines.

- **Consistent Reporting Process** Ensure that problem reporting uses the same system for all external and internal users. This requires a certain level of internal discipline to get all administrators—both in systems and applications support—to utilize the same system. This is important because a problem reported externally may be linked to a problem reported internally, which will have more technical detail. An example is a network-shared drive that continues to be inaccessible, as reported by external users. An internal problem report indicates that the network HUB associated with that device has experienced intermittent failures and is schedule to be replaced. These two sets of information provide a quick level of diagnostic information that should lead the storage administrator to verify the NAS device internally and take steps toward a resolution that may require a different path to the storage.

- **Historical Reporting and Trending** One of the most valuable assets problem tracking systems provide is a database to view the history of a component or device, its past problems, and resolution. This can be extremely important as both SAN and NAS components are tracked and analyzed for consistent failures or weaknesses. For example, historical trends regarding complex switch configurations can be very helpful in identifying weaker components or trends toward failure as port counts become higher and higher and interswitch linking becomes more sophisticated. This proves a valuable tool in anticipating potential problems and thereby scheduling changes prior to a failure. This approach offers a proactive answer to problem management and subsequent downtime statistics.
- **Experience and Expertise** Tackling problems is one of the quickest ways to learn a system or technology. Using problem tracking information and configuration management information provides effective training for IT personnel new to storage networks.

Closing the Loop on Serviceability

Effective serviceability requires the establishment of both configuration management and problem management. Starting with a set of user requirements that result in the capacity plan, followed by the installation, the critical information contained with the physical and logical configuration documentation drives the effective response to problems.

Problem management establishes a method to effectively deal with problem servicing and the proactive servicing of the configuration. Key among problem management activities is the definition of prioritization. Although this can be a volatile activity, it's best kept as an internal data center activity where end-user data utility is a key goal. Problem identification is central to establishing credibility within the data center. Given SAN and NAS configurations, both can be analyzed from an outside-in problem identification perspective. Although SANs demonstrate a much higher level of complexity, they nevertheless lack the necessary tools to enhance the problem identification process. On the other hand, NAS configurations, given their operation as an existing network component, can take advantage of current network tools. NAS also has a level of maturity that offers internal diagnostic information, thus providing efficiencies in problem identification and resolution.

Without an effective practice for servicing the storage network, the complexities and potential confusion within configuration analysis and problem diagnostics will exacerbate system outages. Both configuration management and problem management are critical elements to the cycle of storage network management.