

The  
Complete  
Reference



# Appendix C

**SAN/NAS Management Case  
Study: The Southwestern  
CD Company**

The most important aspect of business application performance is the ability to access data. Data stored on disk, tape, optical media, or temporary cache, using the traditional direct attach storage models will be less susceptible to performance and operational anomalies than data stored through the network model. Although the location of the data should be irrelevant, 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 ability of an IT staff to locate, identify the cause, and correct these anomalies becomes more problematic as storage configurations move into SAN and NAS configurations.

The key to successfully monitoring and managing the performance of any application starts with the consistent and proactive management of the critical paths to application data. In most cases, this is done with an assortment of vendor- and user-developed tools that support management of storage configurations. In Part VI of this book, we discussed many of the accepted management disciplines that become the foundation for a well-run IT organization, none of which will be more important in this case study than performance management. Having said that, performance management is tied to the necessary yet mundane activities of configuration management and the need for an orderly problem management system of processes and procedures.

In this study you will discover that in the storage networking area, the availability of tools and coordinated systems will be lacking greatly. As a general industry trend, management tools lack the implementation of successful technologies by three to five years. This situation is becoming even more significant and problematic as the integration of infrastructures needs to be taken into account in both small and large configurations. As the chapters on connectivity options for both SAN and NAS point out, the effects and counter-effects of storage networks are experienced on an enterprise-wide basis.

Unfortunately, as this case study will demonstrate, today's storage management industry remains in an often misguided and confused startup mode with point products and integrated storage management architectures and standards useful only to those companies selling industry controversies. However, storage management has become an effective catch-up mechanism for large vendors, especially those in the software management sector, who have more integration challenges than the small startups.

Nevertheless, regardless of vendor size, the current assortments of products are characterized by the following deficiencies, to name only the major ones:

- No correlation of functions among distributed components
- No collection of statistically useful data
- No proactive performance, configuration, or capacity trending
- No identification of "root cause" to problems

All this makes for a challenge when you are faced with installing, managing, and maintaining a storage area network and network attached storage solutions, in addition to the challenge the SAN and NAS vendors have in compensating for these inadequacies. This study will illustrate many of the problems and potential solutions when the Southwestern CD Company deals with its initial SAN experiences.

## **The Deferred Responsibility Scenario— or Who’s Doing What to Which Data?**

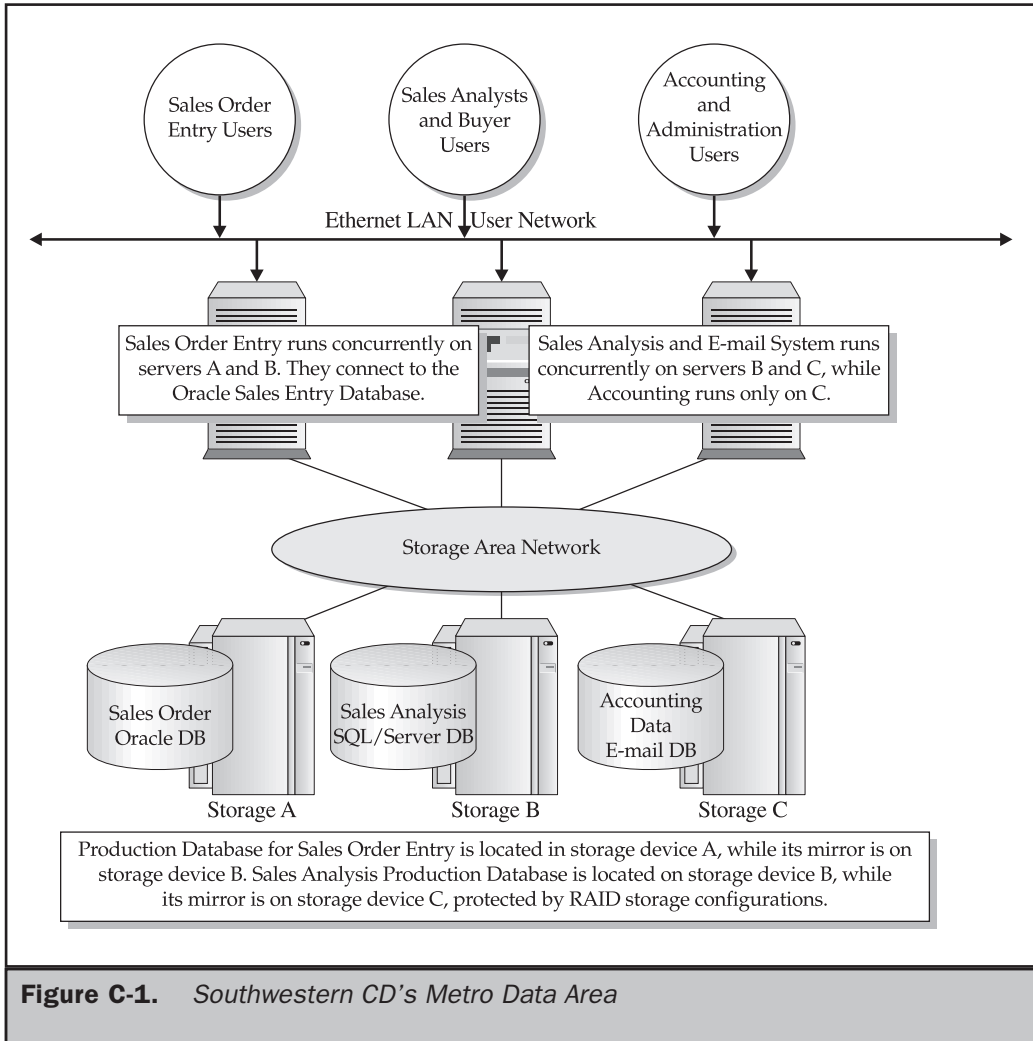
The Southwestern CD Company distributes all types of entertainment CDs to both wholesalers and retail outlets large and small across North America. From its headquarters in Scottsdale, Arizona, the company ships on average 250,000 CDs a week. By purchasing directly from several manufacturers at optimal prices, the company has achieved revenues of approximately \$51 million per annum. Its profit margin is based upon moving product effectively to its retailers to leverage the current entertainment fashion for both music and movies. Because it competes with other, larger distributors, Southwestern must offer better services and a more competitive price structure.

The CD distribution industry is based upon time to market and can rely heavily on technology to facilitate and provide a competitive advantage to the ordering process. CD products are purchased from manufacturers with a return policy that allows the retailers and, subsequently, the distributor to return unsold products. These market dynamics are reflected in the price, generally at the retail level, although the competitive positioning based on availability and amounts can be experienced at the distributor level.

An IT department of 15 people supports the company. Ten permanent personnel and five contract programmers staff the department. They do not have specific storage expertise, but they rely on their systems administrators to handle the operating system infrastructures that include the storage configurations. Their configuration is Microsoft-based with the exception of the database servers, which are UNIX machines running a popular relational database system. They have recently implemented a SAN to consolidate their databases and provide more responsive order-entry and sales-entry applications.

In conjunction with order-entry and sales-entry applications, the company runs an accounting application, a sales analysis application, and corporate e-mail, all from within its Metro Data Area (see Chapter 2). The recent implementation of the SAN has consolidated data from all of these applications within the available storage resources. All applications are now centrally accessible through the SAN.

The company’s configuration, business applications, and client network is depicted in Figure C-1.



**Figure C-1.** *Southwestern CD's Metro Data Area*

## The Problem

The company's new sales-entry application has experienced periodic slowdowns in its performance. In certain instances, the database has had to be restarted and reorganized to resume operation. Although the rationale behind this has yet to be analyzed or understood, restarting appears to restore performance for a time—all of which takes time away from the sales department, which cannot enter sales data while the system is down. Today's scenario, however, is further complicated. It's the end of the month and the day begins with a handful of frustrated phone calls from sales personnel, who

either can't log in to the application or are suffering because the application is running extremely slowly.

## Business Situation

At the end of the month, the sales department really needs to hit its numbers so that the orders can be booked for this quarter. Order quotas come from the Southwestern executives, are part of the quarterly goals, and are the basis for the financial incentives for the sales management staff. They rely heavily on the system to be up and performing optimally to enter as many orders as possible. Over and above is the new pricing structure on the hot new CD products they can sell to their largest retailers, to which they are still trying to sell.

## Technical Scenario

The IT department is short staffed, because the database programmer recently quit to take a job at a bank. The application database administrator (DBA) now has new and additional responsibility for the systems aspects of the production databases. The periodic slowdowns happened before the SAN installation; however, they were never completely debugged and root causes were never identified. The problems appear to have gotten worse as the SAN went into production. All of the user table data has been transferred to the SAN disk arrays, which was the last activity the recently departed database programmer accomplished.

## Problem Dialogue

The following is a summarization of the IT department's events and activities in determining the cause and solution to the problem.

*9:45 a.m.*—After numerous calls to the Help Desk over the past hour, the IT operations manager fills out a “trouble ticket” and assigns it to the DBA.

*10:45 a.m.*—Busy implementing the next release of the database while parsing additional database design changes, the DBA, having worked through most of the night, arrives to work late. The DBA initiates the database management program, which provides performance metrics on the database, leading to the following conclusions: the sales-entry database is up and running, and a number of deferred writes are filling up the temporary storage. The DBA decides the problem is an operating system anomaly and reassigns the “trouble ticket” to the Windows administrator.

*11:00 a.m.*—The Windows administrator, who is also the LAN administrator, is busy implementing a new LAN segment to facilitate the expansion of the accounting department. Consequently, he must be paged when there is no response to the “trouble ticket” within the allotted time.

*11:15 a.m.*—The Windows/LAN administrator logs in to the sales servers and initiates the OS management software, which provides performance metrics on each of these systems. Two anomalies become apparent: client transactions are being queued

at an increasing rate, and this is putting stress on the system as memory and temporary space within the database becomes saturated. System-level paging has increased to an alarming rate. (See Figure C-2.)

*11:30 a.m.*—After a brief discussion with the DBA, the Windows/LAN administrator and DBA concur that a reboot of the sales system and reinitialization of the databases will clear up the problem. By this time, the entire sales-entry system is down. A message is broadcast via e-mail and voice-mail stating that the sales-entry system will be rebooted during lunch. The VP of sales estimates \$500,000 in sales has not been entered. He is not pleased.

*12:30 p.m.*—The DBA and the Windows/LAN administrator discover that although the servers are rebooting, the database is extremely slow coming up. The Windows Performance Monitor tool now indicates that few I/Os are being serviced from the sales servers. After further discussion, the two hit upon the possibility that some external application is impacting the reboot of the databases.

*12:45 p.m.*—Inquiries are made and the two discover a large archive job has been running since 8:00 A.M. Initiated by the Sales Analysis department, the job was to dump two years' of sales history from the SQL Server database running on the ADMIN server. Not surprisingly, the sales history database is stored within the same storage subsystem as the sales-entry database. The DBA and the Windows/LAN Administrator (who've ordered in lunch by this point), attempt to contact Sales Analysis management to get permission to stop the archive job and restart it later that night or over the weekend. The entire Sales Analysis management team is out to lunch, of course.

*2:00 p.m.*—The DBA finally contacts a Sales Analysis manager and negotiates a stop to the archive job, which will restart later that night—but with the mandate that he (the DBA) have the data archived and reorganized for his department by 9 A.M. the next morning. The DBA is not pleased.

*3:00 p.m.*—The Windows/LAN administrator and DBA terminate the archive job and restart the sales servers and sales-entry databases, all of which come up running at optimum levels. The second message is broadcast via e-mail and voice mail: "System's up. All's well that ends well."

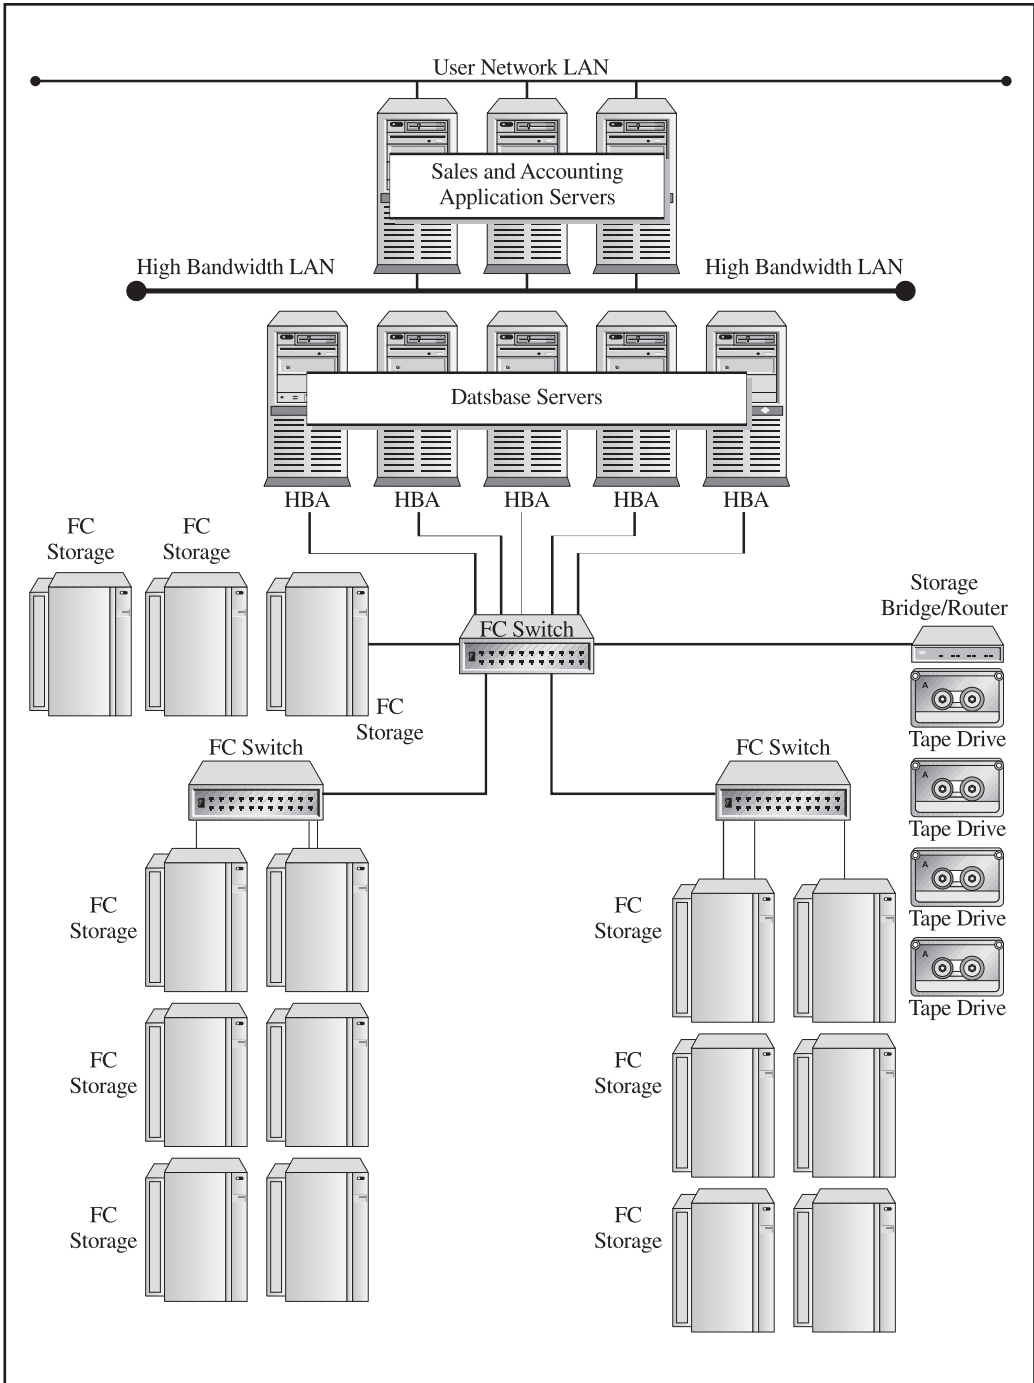
*3:30 p.m.*—Sales personnel are finally back online and able to enter sales data into the system.

## Epilogue

The analyses surrounding activities of problem identification, determination, and correction are indicative of the state of management within storage networking environments. The following epilogue, or post analysis, can provide some insight into the hidden problems that can develop as storage configurations move into a networked setting.

### Root Cause

In this case, unexpected I/O activity from a single server impacted the capacity of the I/O subsystem and storage network to the degree that it severely limited the operation of the sales-entry application and access to their data sharing the same network.



**Figure C-2.** Southwestern CD's Systems Configuration

## Solutions

Unscheduled archive jobs were terminated and rescheduled, sales servers were rebooted, sales-entry databases were reinitialized, and recovery programs were run.

## Downtime

In effect, the sales-entry application was down from 8:30 A.M. to 3:30 P.M., a total of 7 hours. Due to the sales-history database not being archived and reorganized, it was unavailable for a longer term, some 12 hours.

## Effect on Business

As a result, sales was unable to log all of the day's orders, subsequently losing an estimated \$1 million in backlog business. Downtime of the sales-entry database also adversely affected other departments, including purchasing, whose buyers were unable to meet order deadlines to record manufacturers, which will increase costs by approximately \$2 million for the next month.

## Effect on IT

IT management reviews the steps taken to resolve this problem and concludes the following:

- More correlation of information across the systems is needed.
- Greater attention should be given to the "root cause" of a problem.
- No integrated tools are available to coordinate the actions of the Windows/LAN administrator and the Oracle DBA.
- The Windows administrator and the DBA must work toward proactive trending in estimating storage resource requirements.
- Additional storage restrictions must be put in place to prohibit processing intrusion into the Sales production environment.

## Evaluation

This scenario can be evaluated in three ways:

- Business rules and policies
- The processing configuration limits and choke points
- What available tools and practices are currently available that are relevant and appropriate to this scenario

## Business Rules and Policies

The Southwestern CD Company, like many companies, lacked a complete evaluation and definition of the business applications it was supporting. Prioritizing sales entry,

sales analysis, and other workloads would have provided metrics for dealing with many of the timing issues. Defining these rules and policies is the responsibility of the business user and the IT department and should be performed up front, if only on a macro-level (that is, sales-entry must have highest availability of all applications).

## Knowledge of Processing Configurations

Not knowing the processing configuration, especially the storage network configuration, was especially harmful in this case. More importantly, knowing and understanding the limitations of the storage network in context to the supported servers masked the root cause of the problem. Finally, and most importantly, was the lack of association between the storage network and related devices to the business application. Had this association been made, the problem would have been identified within 15 minutes and resolved within 30.

## Available Tools and Practices

Unfortunately, in this case, any available tools are probably of little use in identifying root causes. The following points provide a quick summary of the value of the major tools currently available:

- **Storage resource management (SRM)** Defensive, reactionary information received after the fact. SRM tools are resource-focused, operate without business policy, and provide no correlation of information across distributed servers; in addition, they have no relationship to the business application.
- **Backup/recovery** Adding new and/or bolstering existing backup/recovery software to the storage network. This is complementary, at best, to resolving the problem. The benefit of faster recovery time is not part of the root cause of the outage. Although necessary, backup/recovery bears no relation to the business application.

Although more software tools are becoming available, they continue to provide disparate, incompatible, and inconsistent levels of information on the storage infrastructure. 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, to 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 resources on business applications.

The deficiencies within these 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 only a few of the major deficiencies challenging today's management tools, as well as the inefficiencies surrounding business application availability:

- *No correlation functions among distributed components.* Today's business applications are distributed, which means that events happening on one server can seriously degrade performance throughout the entire system. The ability to correlate important aspects of performance information as it effects the business application is currently unavailable.
- *No proactive trending.* Today's IT managers are expected to drive the bus effectively while monitoring performance through the rearview mirror. Literally all reporting and trending is historical. The information that ultimately reaches the IT user is about incidents that have already occurred and provides little value in determining real-time solutions to business applications that perform poorly.
- *No identification of root cause to problems.* The effects of the conditions stated previously means it is unlikely that the information supplied to the IT user will provide any sort of root cause as to why a business application is down, not to mention any information intended to identify and correct the problem. Most activities address only the symptoms, leading to reoccurring results.

As far as practices go, this case study illustrates a competent and responsive IT department with processes in place for problem notification and assignments. However, the specific expertise and focus of the individuals contributed to the time it took to pinpoint the problem—that is, the database was up but was not processing any transactions.

The following points present alternatives to consider when providing a long-term solution:

- **Reconfiguration/isolation** Adding hardware (servers, switches, and so on)—also known as “throwing money at it.” Although a possible long-term solution, this can essentially convolute the problem and put storage isolation as part of the IT solution matrix, thus limiting the inherent benefits of storage networking. Also, reconfiguration/isolation provides no direct association or long-term benefit to the applications.
- **Increased manpower** Hiring more people to monitor and manage the storage infrastructure by business application—also known as “throwing money *and* people at it.” This could be a solution, but a highly improbable one, given the monitoring requirements, expertise required, and associated costs. This solution does in fact relate to the business application, but only at a prohibitive cost.

## Final Analysis

More powerful servers and storage subsystems are quickly evolving into Metro Data Areas characterized by high data traffic, inevitable congestion, and complex problems

relating to data availability. The Internet is also driving the Metro Data Area theory by increasing exponentially the number of data highways leading in and out of these already congested areas.

Business applications are distinguished from other applications through their direct association to the business, their size, and their mission-critical requirements. They are supported within storage systems by practices and products that manage various aspects of maintaining the storage devices and data. Business applications can be enhanced with storage networking technology, although a host of problems remains that will cause difficulties relating to data availability. Currently, no tools directly relate the business application to the storage network.

Through a composite IT scenario, the realities behind operating with a storage networked environment were illustrated in this case, as well as the challenges in problem identification, isolation, and solution within these configurations. Although a SAN configuration was used, a NAS configuration could just as easily be substituted for the same effect. Evaluation of the scenario indicated that best practices and current tools would not minimize or stop these types of availability problems.

## Steps to Better Storage Networking Management

As demonstrated in this case study, the association of business applications to computing resources remains for the most part at macro-levels, at best, and not at all within new technologies, such as storage networks. Storage networks, given their disparate components and critical interrelationships, have evolved into an infrastructure category within the data center. As such, they have become a tremendous challenge to manage as a complete entity, much less integrated into the logical set of applications and supported systems. The following provides additional information on what's required for new tools. This is further augmented by best practices for managing storage networks.

### The Need for Next-Generation Performance Management Tools

The most critical element of performance for a business application is its availability to its own data. Consequently, a level of performance tools is needed that closely associates the business application with storage networking resources. These tools should provide a level of functionality to perform the following:

- Correlate business application performance information across servers and storage networks
- Provide metrics and monitoring for proactive trending of performance and take appropriate action based upon previously established business rules
- Identify root cause areas, if not components, to IT support personnel and make appropriate recommendations based upon previously established business rules

Such a set of tools would enhance IT management's ability to manage business applications within increasingly complex application and storage networked configurations. These tools would move IT professionals closer to the critical elements of business application performance.

Storage networking best practices to observe include the following.

- *Eliminate unnecessary downtime.* Decrease the negative business impact of storage networking and enhance the inherent abilities of storage networking technologies.
- *Automate performance management tasks.* Decrease the cost to IT management and staff, and maintain specialization in storage networking as implementation and production use increases.
- *Learn/associate business to infrastructure.* Root cause analysis, proactive processing, and management can provide solutions before outages occur.

The following guidelines can provide some visibility into dealing with the Southwestern CD company scenario:

- *Focus maintenance efforts on critical resources—before they go inoperable.* Identify the critical resources through effective configuration management activities and processes. These can help in focusing efforts on current configuration and processing scenarios.
- *Proactively manage growth and change within the storage networked configuration.* As the storage network moves into production, be proactive through an effective problem management and ongoing capacity plan. This will allow the identification of problems and changes to capacity. Further accommodations can be helpful through an effective change management program that will identify and focus on processing anomalies and activities.
- *Gather value oriented performance data.* This will probably be the most intrusive and time-consuming task, given the lack of availability and the programming necessary to gather information from disparate sources such as MIBS, Common Information Model, and vendor statistical collections. This can also provide valuable input to vendors and industry during problem-resolution scenarios and activities.