MANAGING HETEROGENEOUS STORAGE AREA NETWORKS

By Robert Wright, President and COO
SANavigator Inc.
San Jose, Calif.

Storage Area Networks (SANs) offer many advantages to companies attempting to meet the constantly changing data storage requirements of today's large, rapidly growing networks. SANs provide high-speed, scalable, networked storage solutions that separate the storage devices and storage access from the enterprise network. This provides the opportunity to centralize storage management and to relieve storage-related bottlenecks in the LAN and within servers. SANs can solve the problem of any-to-any connectivity between storage devices and servers. A SAN also offers new approaches to storage, such as disk and tape pooling, heterogeneous data sharing and off-the-network/serverless backup and restore.

While a SAN can significantly improve data and storage device management, its actual performance depends on a number of factors. The first of these is the quality, features and interoperability of the storage connections and devices (e.g. Fibre Channel host bus adapters (HBAs), hubs, switches, bridges, routers and cables) that are employed in the SAN. With inferior components it can be very difficult to maintain data availability. With such devices it is often necessary to spend an inordinate amount of time trouble-shooting problems, re-routing data paths and installing components. It is also unlikely that one manufacturer will be able to provide the best features and quality for every SAN hardware and software component. This makes it highly desirable to choose the best available components for each function within the SAN.

Just as important as the quality and feature set of the devices is how effective the SAN management application is in tying these devices together and simplifying the native complexity of the storage network. An IT manager needs to configure, monitor, and maintain the SAN, as well as plan for, and accommodate growth. Without a robust management application designed to meet the manager's requirements, even the best SAN installation can represent an administrative nightmare.

The Challenge of Managing "Best of Breed" Devices

To ensure that each component in a SAN provides the desired functionality and quality, it makes sense to select "best of breed" components. What represents the best product for one installation may not always be the best choice for another. Requirements may vary significantly, for example, in the areas of availability, capacity, performance and cost. A SAN that consists of "best of breed" switches, hubs, routers, storage subsystems, cables, connectors and software is therefore most likely a heterogeneous, multi-vendor network.

Because it is common to add components to SANs as specific needs arise, it is also typical to have a heterogeneous network whose components are of varying quality and capability. For economic

reasons, it may not be desirable to replace older components in a SAN until they fail or become obsolete, but it is still necessary to manage the network.

A heterogeneous SAN based on "best of breed" devices and a mixture of older and newer components can present significant interoperability and management problems. That is because the storage industry lacks widely adhered-to standards for devices and for the software that manages them. While some standards, such as the Fibre Alliance MIB (Management Information Base) are emerging, these standards are not fully implemented by many of the leading hardware vendors. Instead, vendors often employ their own, proprietary MIB's and API's (Application Programming Interfaces).

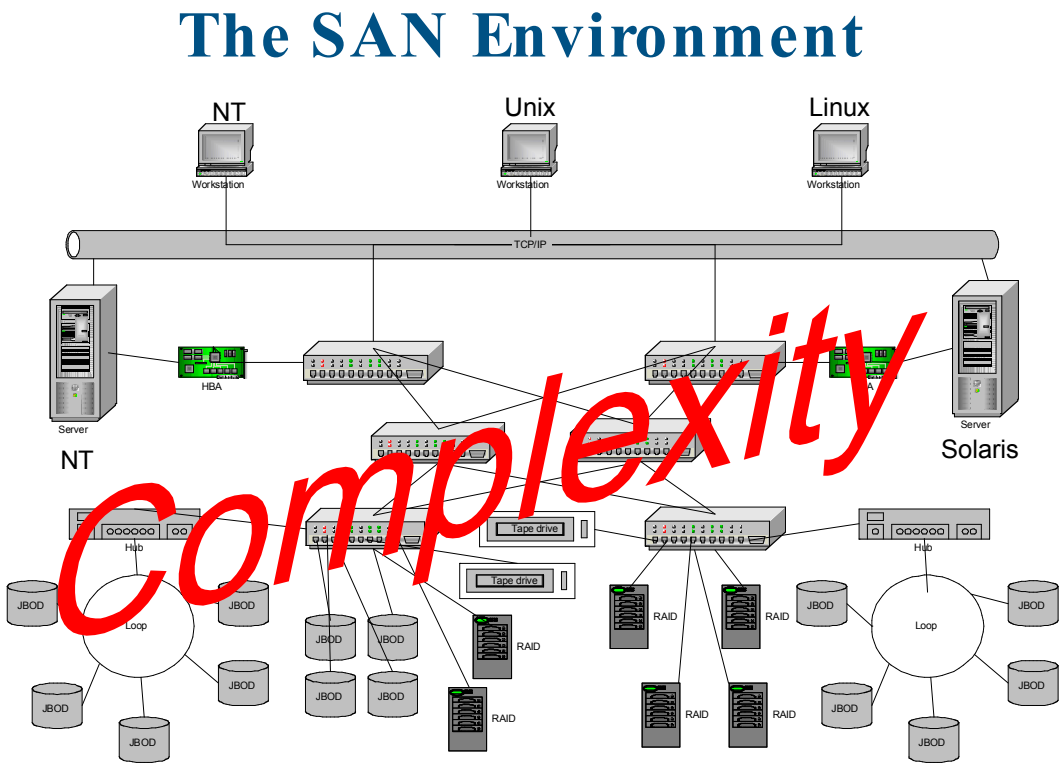

Figure 1: In a highly complex SAN, such as the one illustrated here, there can be a wealth of devices from different vendors, multiple operating systems and innumerable connections to track. Having software that is capable of managing this diverse system is necessary.

Addressing the Challenge

How is it possible to avoid the interoperability and management problems that can occur in a heterogeneous (multi-vendor) SAN? The best way is to make device interoperability a consideration during the component purchase cycle, and to select a SAN management application

that offers a great deal of flexibility and robustness in working with the current and forthcoming standards, features, *and* the proprietary interfaces of all the major equipment and software vendors.

Component Selection

Standardizing on one best-of-breed vendor for each *type* of SAN component can help to minimize the complexity of heterogeneous SANs (assuming a vendor ensures interoperability among its own, similar products!). For example, standardizing on one switch vendor for each SAN installation provides the ability to cross-connect the devices for improved availability in the event that one switch fails. Even if the SAN installation includes only one vendor per component type, the interoperability between the various types of components in the SAN must be addressed. Vendors who support emerging standards should receive preference over those who don't. In addition, most leading vendors have (or subscribe to) interoperability certification programs, through which product interoperability is extensively tested. Ensuring that the selected SAN components interoperate effectively in these labs through the applicable certification programs will reduce the likelihood of device interoperability problems.

Selecting the Right SAN Management Application

Once the SAN devices are chosen, selection of the right SAN management tool is key. This application should be flexible, portable, scalable and reliable. In addition, it should be simple to use and protect the SAN from novice administration errors.

Application Flexibility

To be truly flexible, SAN management applications should offer some type of plug-in architecture, so that new devices can quickly be incorporated into the tool. The software should also be architected so that it can easily support the addition of new discovery protocols (e.g. Jiro™). Additionally, the software should have a layered, modular architecture, so that it is possible to quickly and easily integrate any feature in the SAN management application with the APIs of other vendors' products. A layered architecture can also reduce the amount of code that any modification (e.g. feature addition) impacts, because functions are contained within specific modules rather than distributed throughout the application. This greatly simplifies software maintenance, and reduces the chance of a code modification introducing new problems in the application. Archaic (monolithic) software architectures that do not provide this flexibility increase the chance of experiencing debilitating application "bugs," and interoperability "issues." Additionally, poor software architectures may significantly increase the number and frequency of delays in receiving support for new features, devices, operating systems and protocols.

Each of the hubs, switches, tape drives and other devices in a SAN may have its own management software. To keep from having to monitor hundreds or thousands of consoles, all of these devices should be managed from a single console. State-of-the-art SAN software management applications can manage the SAN on-site or remotely, providing management access from anywhere.

To simplify SAN management in a heterogeneous environment, the SAN management application should make it easy to operate software tools from a variety of vendors. Automatically launching native software tools for any of the devices (or management functions) in the SAN makes possible unique device vendor error reporting and correction tools, and allows the software to scale with added devices and features.

When a manager wishes to add a device from a new vendor, the SAN management application should allow association of the new vendor's devices with the desired software. A simple launching mechanism should enable the tool to find and launch the unique vendor application.

Storage management applications should provide integration with enterprise level network management applications, in case these applications are someday added or merged with the other administration tools. This should include, as a minimum, pass-through of SNMP traps (events). Advanced management applications can also forward events to multiple "listeners" (e.g. for remote or alternate location monitoring), provide topology import/export, and bill of materials export in standard formats.

Portability

The management software for heterogeneous SANs should be portable to a variety of operating systems. Enterprise networks that utilize SANs often include Windows NT, one or more versions of Unix, and the Linux operating system in various departments within the enterprise. SANs often originate with the Unix-based storage requirements of development groups, and later expand to include NT-based administrative storage. A Java-based, enterprise-wide storage and SAN management platform that supports remote usage capabilities, provides the easy portability that a SAN requires. To provide outstanding portability, it is optimal for the SAN management application to be 100% Java.

Scalability

SAN management software must be able to scale with the growth in devices and storage requirements. Scalability is exceptionally important, since storage requirements in a typical enterprise are growing at more than 50% a year. Leading tools allow for dynamic discovery, i.e. SAN components can be added and deleted at the will of the administrator, without having to re-start the SAN. Most vendors also supply a mini-map, or alternative method (e.g. zoom levels) for navigating through large topologies.

The application should also provide a hierarchical user interface to ensure scalability as the SAN grows. If the abstraction levels are not built into the software from the beginning, problems with scalability are probable later on. It is also important that the software vendor "own" the user interface so that it can be optimized for storage networking. To hasten development schedules, some vendors utilize generic third party applications for their user interface. Generic interfaces are not optimized for SAN management, and can result in scalability, configuration, interpretation and support problems when the SAN grows beyond the generic software's capability.

Reliability

The SAN management application should support, as a minimum, out-of-band discovery. If the SAN experiences a comprehensive catastrophic failure, remote out-of-band discovery and monitoring can ensure that the management application is still able to notify the administrator of the failure. If the application only supports in-band discovery, then the application could be disabled when a major SAN problem occurs.

To help ensure a reliable SAN management application, the application vendor should have a well-planned software architecture, and a dedicated, state-of-the-art software quality engineering program. Although beyond the scope of this article, an object-oriented, iterative or adaptive-style type of development program that also involves frequent integration and constant test, including an extensive cross-vendor interoperability test program, are some of the indicators of an excellent quality-engineering program.

Simplicity

The appearance and ease of use of the SAN management application user interface can have a significant effect on the way a SAN is managed. A well-designed interface that is intuitive and easy to use can greatly simplify advanced management functions for the administrator and provide a vital link between the storage network and other enterprise management tools. For example, each switch and router vendor has their own method of constructing server storage zones. The user interface can simplify zoning by providing an intuitive, consistent, and visual method of constructing zones, while relegating the complexity of the different zoning methodologies to the lower (and user-transparent) levels of software.

It is essential that the interface be custom written with SAN management in mind. Generic software, not dedicated to SAN support, does not provide the most helpful visual representations of the devices in the SAN, the way they relate to each other and the way the SAN relates to the enterprise. For example, in generic SAN software, icons for storage devices might look like boxes with an X in them, instead of storage devices.

The network topology should be displayed on the screen in a way that makes sense to someone who is managing a SAN. For example, it helps if the LAN devices are displayed in a different color and grouping than the SAN devices. For easy identification, errors should also be displayed in color, and categorized based on their severity.

Safety

Safety goes beyond the protection of data. It should include protection of the network itself. For example, the SAN management application should allow the display of all information that is discovered, but restrict the user's ability to change those items that can negatively impact the network. By implementing role-based security, a user can only perform those functions that they

are qualified to perform. Advanced configuration and management functions are therefore restricted to the advanced administrator(s).

II. Managing the Heterogeneous SAN

Given a SAN management application that fulfills the requirements for flexibility, portability, scalability, reliability, simplicity and safety, what should the application do for the user? To tie the diverse SAN devices together in a meaningful way, at a minimum the SAN management application must offer:

Automatic device discovery
Topology mapping
Persistence
Status monitoring
Event logging and notification

To be a complete management package, SAN management software needs to manage advanced functions, such as backup, zoning and storage pooling (virtualization). The storage management application should provide these services either directly, or through API's to other applications/devices.

Automatic device discovery, Topology mapping, and Persistence

Organizations often need to support multiple subnets. A SAN management application must be able to automatically discover devices on any subnet in the world, through local or remote monitoring. It also should be able to display multiple subnets on a single management console. Given the subnet address, the tool should be able to support the added subnet immediately, without complex configuration tasks.

After obtaining the subnet address, the SAN Management Software should automatically discover all the devices in the SAN—hubs, switches, HBAs, routers, etc., and their connections. The application should also identify the device vendors, worldwide names and other properties, if available from the device software. If the device is not an identified component, advanced management applications provide the capability for a user-defined device, making it possible to create a complete SAN component map. The characteristics of each component should be available by clicking on the component icon, and should be saved, so that the user only enters information once.

The software must also be able to map the device's location in the network- its place in the network topology. The topology information should be presented in the form of an intuitive visual map that includes all the components, devices and their interconnections, showing groupings of the hosts (or HBAs) and switches.

To provide users with an instant and intuitive understanding of the network, redundancies should be displayed graphically. In the past, this type of visualization, or diagramming of the SAN topology was an arduous task that needed to be carried out with paper and pencil or rudimentary drawing software.

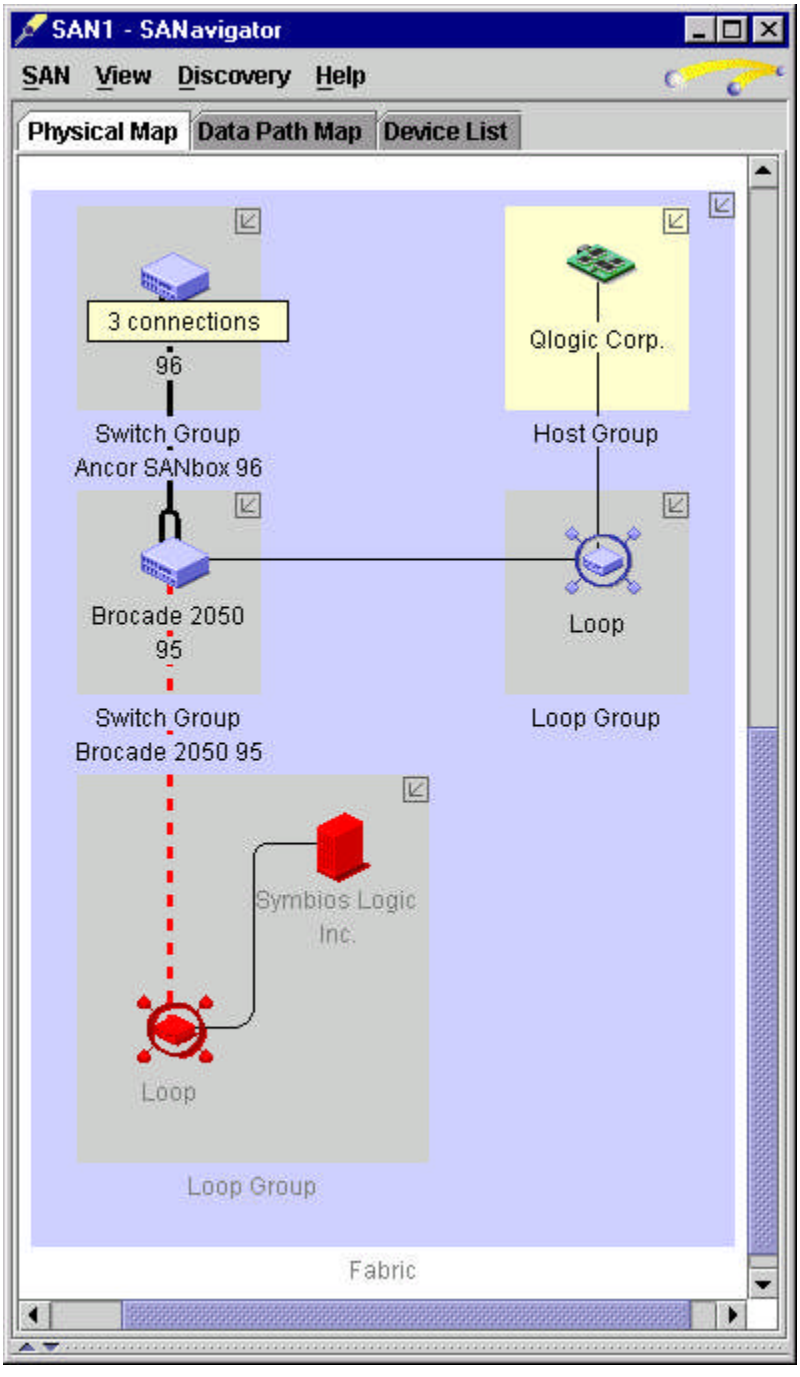

Figure 2: SANavigator, from SANavigator Inc., displays dashed red lines and red icons for failed devices or cables. Redundant connections are shown as a fork at the device connector, and are also displayed numerically during mouse-over. The icons are designed specifically for SAN devices.

When a SAN device is added, moved, or removed, the software must be able to detect it, and adjust the topology display accordingly. Devices and connections that are removed or failed should be displayed differently (e.g. in red) to allow quick identification of trouble spots in the SAN.

The discovery and mapping process becomes more complicated when a device utilizes a vendor's private discovery protocol. The software must still be able to detect these devices, and add them to the topology.

As the SAN management application discovers and maps the devices, it should save the information so that it is "aware" of any changes that occur in the SAN. This feature is called persistence. All user-entered information should also be saved. Persistence, combined with import/export capability, allows the user great flexibility in utilizing the software to identify changes that occur in the SAN (e.g. between administrator shifts), and to re-create saved configurations (e.g. if a SAN needs to be quickly re-established after a company move or local catastrophe).

Status monitoring, Event logging and notification

Management applications should monitor the SAN for changes in device and connection health. Out-of-band discovery via SNMP features the ability to "listen" for device traps (events), such as a redundant power supply failure or environmental problem. If a catastrophic event occurs, however, such as a device failure or a cut/broken cable, the component may not be capable of sending a call for help. The management application should therefore poll the devices every so often just to make sure everything is still functional. Ideally, polling intervals and timeout periods should be changeable by the user. Together, event listening and polling can ascertain the overall health of the entire SAN.

Monitoring services are becoming increasingly important for SANs. There are a growing number of organizations that monitor other organization's SANs, or that have co-located storage around the world. The SAN management application should therefore provide remote monitoring capability, including an event forwarding service and remote notification of any failure. All events should also be logged, so that a record may be kept of SAN health and performance. More advanced applications allow selected event types to be aggregated at user-specified time periods, and sent via SMTP to pagers or e-mail locations.

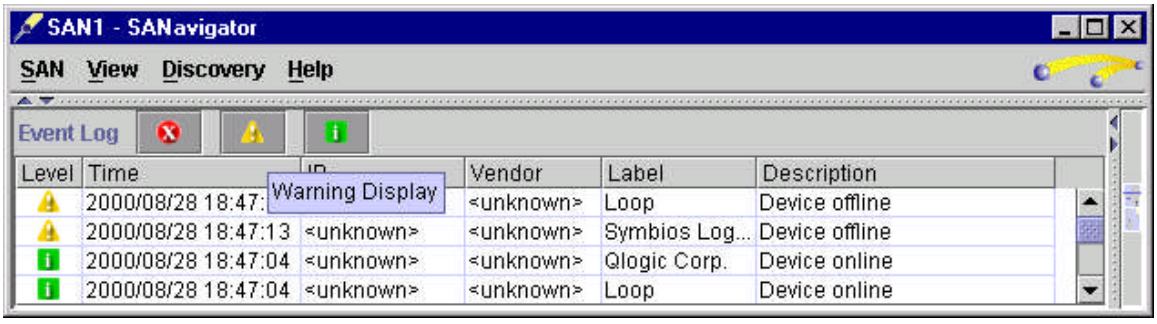

Figure 3: SANavigator, from SANavigator Inc., displays event information, separated into three severity levels.

Advanced Functions (e.g. Backup, Zoning and Storage Pooling)

Advanced function support should be provided in the software itself, or through an ability to launch other software modules. An attraction of utilizing the management software to launch some functions from a particular device is that the user can create virtual appliances. For example, a 3rd party backup application might be associated with a particular host, or even a NAS box on the front end of the SAN, and the launching host could be renamed "Backup Appliance".

One of the major advantages of a SAN is that the storage devices and data paths can be organized in many different ways, to effectively manage the transport and storage of data. One of the most important configuration issues to address is data path zoning. This is primarily a concern if more than one operating system exists in the SAN. To prevent one rogue operating system (OS) from "taking over" all the available storage, the storage can be divided, or zoned for each OS. Additionally, the user may wish to segregate sensitive data, so that the data can only be accessed by a particular host. Because the leading vendors vary so much on their implementation of zoning, it is advantageous if the management software provides this service.

The SAN management software that supports data path zoning should provide an intuitive interface that allows the manager to select both the host and storage ends of the data path. For example, the software might allow the manager to highlight the particular path of interest between two devices, and enable or disable it.

An alternative to zoning for multiple operating systems is to pool the storage. Utilizing universal file system software or specialized SAN controllers, the administrator can enable multiples OS's to share data. The management application should have the flexibility to allow the user to implement storage pooling and zoning methodologies through a simple user interface.

What Features Should SAN Management Software Include?

Feature	**Key Benefit**
Plug-in architecture	Timely support of new devices
Support for the addition of new discovery protocols	Timely support of new protocols
Layered architecture	Can handle proprietary management systems

Single management console	Ability to monitor/manage the entire SAN with a single interface
Automatic launch of native software tools	Unique device vendor error reporting and correction
Association of devices with the desired subsystem software.	Flexibility in application launching
Pass-through of SNMP traps (events)	Interoperability with other applications
Forwarding of events to multiple listeners	Remote monitoring
Topology import/export	Ability to quickly restore a SAN after a move
Bill of materials export	Asset management
Multiple OS support (Java-based)	Portability to a variety of operating systems
Hierarchical user interface	Scalability
Single vendor owns the management software and the user interface	Optimizes scalability, configuration, presentation and support
State-of-the-art software quality engineering program	Ensure reliability
Intuitive management interface	Easier to visualize components of SAN. Positive impact on administrator productivity.
Restricted ability to change those items that can negatively impact the network	Protect the network from serious administration errors
Role-based security	User rights reflect authority and capability
Automatic device discovery	Enables SAN monitoring
Topology mapping	Illustrates all SAN connections and devices
Persistence of topology and properties	Enables software to notify administrator of topology changes
User-specified properties	Enables user to track properties not reported by the devices
User-specified polling interval	Allows optimization based on user needs and network performance
User-specified polling time-out	Allows optimization based on user needs and network performance
Status monitoring – polling	Allows application to identify catastrophic component events
Status monitoring – event listener	Allows application to identify non-catastrophic component events
Event logging	Provides a historical record of events
Event notification (e-mail and pager support)	Eliminates need to constantly monitor the SAN
Mini-map navigational aid	Aids navigation in large topologies
Data path zoning	Change data path zones
Discovery and mapping of subnets	Support subnets anywhere in world
Remote Monitoring services	Enables remote monitoring of SAN

Conclusion

The management of heterogeneous SANs requires due diligence in component selection, including both hardware and software. Selecting best-of-breed hardware and software that can handle such a complex environment will greatly simplify SAN management, while providing enterprise-wide data protection.

Robert Wright *is the president and COO of SANavigator, Inc. (www.sanavigator.com), in San Jose, Calif.*