



SNIA

Storage Networking Industry Association

Managing a Fibre Channel Storage Area Network

Storage Network Management Working Group for
Fibre Channel (SNMWG-FC)

November 20, 1998

Editor: Steven Wilson

Abstract

This white paper describes the typical Fibre Channel Storage Area Network (SAN) environment from the management perspective, and identifies the challenges of managing that environment. As a solution to these challenges, a SAN management architecture is proposed and a roadmap is provided that outlines the tasks and deliverables of the Fibre Channel Storage Network Management Working Group

Introduction

The emergence of storage area networks (SANs) has created the need for new storage management tools and capabilities. While SANs provide many benefits such as lower cost of ownership and increased configuration flexibility, SANs are more complex than traditional storage environments. This inherent complexity associated with storage area networks creates new storage management challenges.

The prominent technology for implementing storage area networks is Fibre Channel. Fibre Channel technology offers a variety of topologies and capabilities for interconnecting storage devices, subsystems, and server systems. These varying topologies and capabilities allow storage area networks to be designed and implemented that range from simple to complex configurations. Due to the potential complexity and diverse configurations of the Fibre Channel SAN environment, new management services, policies, and capabilities need to be identified and addressed.

This white paper assumes that the reader is familiar with the basic concepts of systems management. However, for the reader's benefit, a review of basic systems management concepts is provided at the end of this paper.

The audience for this white paper includes:

- Systems Management Vendors
- Fibre Channel Hardware and Software Vendors;
- System Integrators;
- System, Storage, and Network Architects.

The SNIA would like to acknowledge all members of the Storage Network Management Working Group for their contributions to this white paper. In particular it would like to thank the following people for providing content: Harry Aine, Kim Banker, Kevin Collins, Kumar Malavalli, Michael O' Donnell, Brian Reed, Doug Shue, Richard Taborek, Andrea Westerinen, and Steven Wilson (who also edited this document).

The Fibre Channel SAN Environment

Historically in storage environments, physical interfaces to storage consisted of parallel SCSI channels supporting a small number of SCSI devices. With Fibre Channel, the technology provides a means to implement robust storage area networks that may consist of 100's of devices. Fibre Channel storage area networks yield a capability that supports high bandwidth storage traffic on the order of 100 MB/s, and enhancements to the Fibre Channel standard will support even higher bandwidth in the near future.

Depending on the implementation, several different components can be used to build a Fibre Channel storage area network. The Fibre Channel SAN consists of components such as storage subsystems, storage devices, and server systems that are attached to a Fibre Channel network using Fibre Channel adapters. Fibre Channel networks in turn may be composed of many different types of interconnect entities. Examples of interconnect entities are switches, hubs, and bridges.

Different types of interconnect entities allow Fibre Channel networks to be built of varying scale. In smaller SAN environments, Fibre Channel arbitrated loop topologies employ hub and bridge products. As SANs increase in size and complexity to address flexibility and availability, Fibre Channel switches may be introduced. Each of the components that compose a Fibre Channel SAN must provide an individual management capability, and participate in an often complex management environment.

Due to the varying scale of SAN implementations described above, it is useful to view a SAN from both a physical and logical standpoint. The physical view allows the physical components of a SAN to be identified and the associated physical topology between them to be understood. Similarly, the logical view allows the relationships and associations between SAN entities to be identified and understood.

Physical View

From a physical standpoint, a SAN environment typically consists of four major classes of components. These four classes are:

- End-user platforms such as desktops and/or thin clients;
- Server systems;
- Storage devices and storage subsystems;
- Interconnect entities.

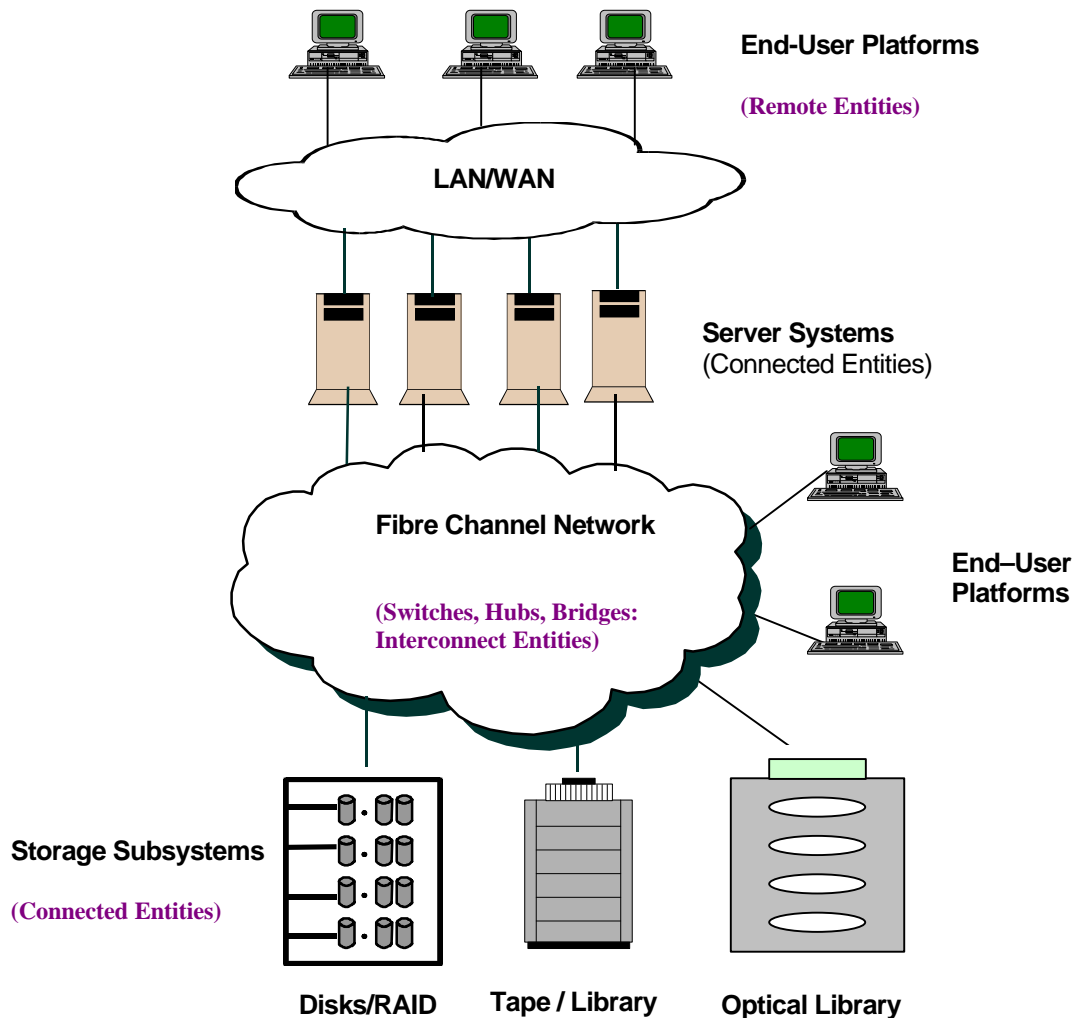
Typically, network facilities based on traditional LAN and WAN technology provide connectivity between end-user platforms and server system components. However in some cases, end-user platforms may be attached to the Fibre Channel network and may access storage devices directly.

Server system components in a SAN environment can exist independently or as a cluster. As processing requirements continue to increase, computing

clusters are becoming more prevalent. A cluster is defined as a group of independent computers managed as a single system for higher availability, easier manageability, and greater scalability. Server system components are interconnected using specialized cluster interconnects or open clustering technologies such as the Fibre Channel - Virtual Interface mapping.

Storage subsystems are connected to server systems, to end-user platforms, and to each other using the facilities of a Fibre Channel network. The Fibre Channel network is made up of various interconnect entities that may include switches, hubs, and bridges.

The figure below depicts a typical physical Fibre Channel SAN environment.



Logical View

From a logical perspective, a SAN environment consists of SAN components and resources, as well as their relationships, dependencies and other associations. Relationships, dependencies, and associations between SAN components are not necessarily constrained by physical connectivity. For

example, a SAN relationship may be established between a client and a group of storage devices that are not physically co-located.

Logical relationships play a key role in the management of SAN environments. Some key relationships in the SAN environment are identified below:

- Storage subsystems and interconnect entities;
- Between storage subsystems;
- Server systems and storage subsystems (including adapters);
- Server systems and end-user components;
- Storage and end-user components;
- Between server systems.

As a specific example, one type of relationship is the concept of a logical entity group. In this case, server system components and storage components are logically classified as *connected* components because they are both attached to the Fibre Channel network. A logical entity group forms a private virtual network or zone within the SAN environment with a specific set of connected entities as members. Communication within each zone is restricted to its members.

In another example, where a Fibre Channel network is implemented using a switched fabric, the Fibre Channel network may further still be broken down into logically independent sections called *sub-fabrics* for each possible combination of data rate and class of service. Sub-fabrics are again divided into *regions* and *extended-regions* based on compatible service parameters. Regions and extended regions can also be divided into partitions called *zones* for administrative purposes.

Fibre Channel SAN Management Challenges

Now that the complexity and scope of the Fibre Channel SAN environment is better understood, the challenges of managing that environment can be considered.

The basic challenge stems from the fact that Fibre Channel SANs utilize a complex network for interconnecting storage devices and server systems. This network is potentially made up of multiple components that have both physical and logical relationships to one another. A breakdown in one component or link in the SAN may manifest itself differently depending on the component that recognizes the condition. However, since a SAN is a network, we can benefit by discussing the management challenges in terms of the basic systems management disciplines.

From a fault management perspective, tools and mechanisms must exist that allow data from multiple SAN components to be correlated and processed. Without the means to correlate data in a SAN, problem isolation is extremely difficult if not impossible. For example, if multiple link failures were to simultaneously occur in a Fibre Channel network, these would need to be correlated with the multiple I/O failures seen in processing a host's I/O requests.

Probably the greatest challenge involves the configuration management of SANs. Due to the large number of components, and the multitude of physical and logical relationships to one another, robust configuration capabilities must be provided for the Fibre Channel SAN. There will be tradeoffs between keeping track of configuration information within the Fibre Channel network, or forcing a central management platform to ascertain Fibre Channel network topology using management mechanisms such as configuration files, name services, and SNMP interactions.

There are also challenges associated with performance management of SANs. Performance information must be provided at a component level as well as an overall system level. Tools and capabilities must exist that again correlate data from a variety of components to provide a system level view of the overall SAN's performance.

Common capabilities must be provided to allow software and firmware updates to be managed from a central management station. These capabilities should allow a generic mechanism to transport updates to components from different vendors.

There are also challenges with the support of accounting or asset management. Capabilities such as standardized SAN resource identifiers containing asset information must be defined. In addition, a common mechanism must be provided to obtain the asset information from SAN resources.

Cluster management is an inherent requirement of SAN management and imposes its own set of challenges. Managing SAN clusters calls for robust solutions for workload management and for single point of control of the cluster's component systems. Data and storage management must be supported within and across clusters. Creating a scalable, single operational management view for SAN clusters enhances productivity and application availability by enabling the automatic detection, monitoring, and control of key cluster components.

Security management in the SAN also creates challenges. For example, restrictive mechanisms such as zoning can have both beneficial and negative consequences for central management. While zoning provides a means to isolate certain storage devices, a centralized manager may need to be a member of all zones and must therefore be a "trusted" entity. The overall security of the SAN then becomes only as secure as the trusted entity.

Ultimately the purpose of providing robust management capabilities for the Fibre Channel SAN environment is to support policy management. Policy management is becoming increasingly important for the enterprise where Fibre Channel SANs will most likely be implemented. In the SAN environment, management capabilities must support the invocation of services chained to well defined policy. Examples of services that require policy are:

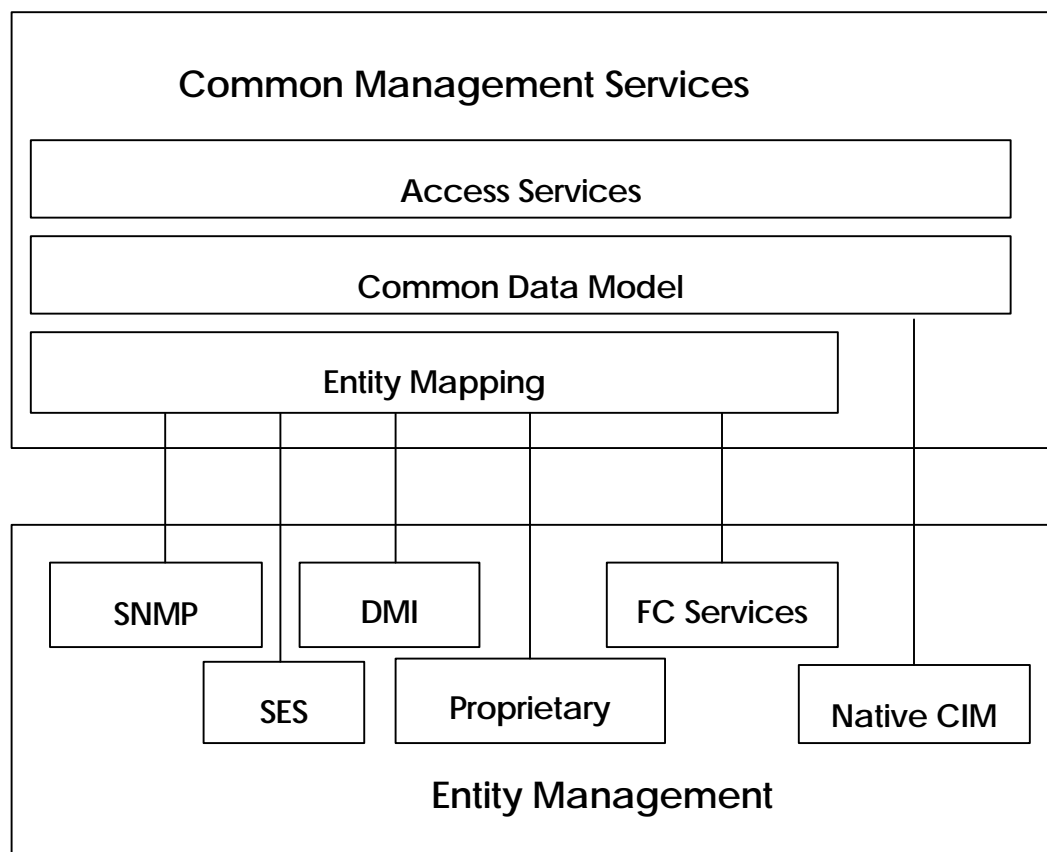
- Disaster recovery;
- Back-up;
- Hierarchical Storage Management (HSM);
- Archival.

Storage Area Network Management Architecture

The primary goal of the SAN management architecture is to provide a common and consistent view of managed resources so that the challenges of SAN management can be addressed. Ideally this is accomplished while preserving existing management implementations for SAN components. In order to achieve this goal, two layers are defined by the Storage Area Network Management architecture. The two layers are:

- Entity Management Services;
- Common Management Services

The figure below depicts a high level view of the Storage Area Network Management Architecture.



Entity Management Services

The Entity Management Services (EMS) layer comprises existing SAN management services and near-term SAN management services. Examples of these services are:

- Simple Network Management Protocol (SNMP) based services;
- Desktop Management Interface (DMI) based services;
- Fibre Channel Services;
- Storage Enclosure Services (SES);
- Proprietary Management Services;
- Native Common Information Model (CIM) Protocols and Services.

The EMS layer accommodates existing management instrumentation and allows additional entity instrumentation to be defined. The services provided by EMS are essentially the super-set of the services associated with each individual entity management service. These services are manifested as either protocol interactions or application programming interfaces.

For example, SNMP management requests may be handled using standard SNMP protocol interactions between an agent on a Fibre Channel switch and the Common Management Services layer. Another implementation could issue SNMP requests through an API provided on a management station.

Common Management Services

The Common Management Services (CMS) layer is the heart of the SAN management architecture. It is based on an object-oriented design model with supporting and associated methods and services. Managed resources associated with SANs are abstracted as management objects by CMS. Management operations for each managed object are exposed through methods that allow attributes and properties associated with managed objects to be monitored and modified.

Management operations effected upon objects at the CMS layer may result in multiple operations being performed at the EMS layer. For example, a method associated with a managed object could be implemented as a combination of SNMP, DMI, and Fibre Channel service operations being performed at the EMS layer.

The Common Management Services layer is divided into three components:

- Entity Mapping;
- Common Data Model;
- Access Services.

The Entity Mapping component translates from the syntax and semantics of the Common Data Model to entity-specific implementations. This component is responsible for interacting with the appropriate entity management services to fulfill a management request invoked by a management application. Entity mapping responsibilities include making the appropriate operation requests to the entity management resources, the collection of responses to those requests, and the subsequent building of the operation response. Entity mapping also collects unsolicited notifications such as alerts, events, and traps from managed resources on the SAN, and packages them for input to the Common Data Model component.

The *Common Data Model* component provides a well-defined and broad view of all managed SAN resources. Fibre Channel SAN resources include adapters, switches, hubs, bridges, storage subsystems, storage devices, and server systems. The definition of SAN management data is complex and spans the representation of simple objects and their properties to complex topological and dependency relationships.

The object design of the common data model supports the following capabilities:

1. **Abstraction and classification** – To reduce the complexity of the problem domain, high level and fundamental concepts (the “objects” of the management domain) are defined. These objects are then grouped into types (“classes”) by identifying common characteristics and features (properties), relationships (associations), and behavior (methods).
2. **Object inheritance** – Sub-classing from the high level and fundamental objects, additional detail can be provided. A subclass “inherits” all the information (properties, methods and associations) defined for its higher level objects. Subclasses are created to put the right level of detail and complexity at the right level in the model. This can be visualized as a triangle – where the top of the triangle is a “fundamental” object, and more detail and more classes are defined as you move closer to the base.
3. **Ability to depict dependencies, component and connection associations** – Relationships between objects are extremely powerful concepts. Before object-oriented management models, management standards captured relationships in multi-dimensional arrays or cross-referenced data tables. The object paradigm offers a more elegant approach in that relationships and associations are directly modeled. In addition, the way that relationships are named and defined describe the semantics of the object associations. Further semantics and information can be provided in properties (specifying common characteristics and features) of the associations.
4. **Standard, inheritable methods** – The ability to define standard object behavior (methods) is another form of abstraction. Bundling standard methods with an object’s data is encapsulation. Imagine the flexibility and possibilities of a standard able to invoke a “Reset” method against a hung device, or a “Reboot” method against a hung computer system regardless of the hardware, operating system or device.

The Common Data Model component may be realized using the Desktop Management Task Force’s (DMTF) Common Information Model (CIM) plus SAN extensions. CIM is a data model that provides a conceptual view of real-world, managed entities (like storage, networks, and systems), which also represents the users, organizations and applications that interact with these entities. CIM attempts to unify and extend the existing instrumentation and management using object-oriented constructs and design. The Common Data Model can build on the object and relationship definitions in CIM to represent the managed Fibre Channel SAN environment.

The Access Service component provides the mechanism that management applications use to interact with Common Management Services. The Access Services represent the exposed object interfaces defined by the Common Data Model and may be realized as XML, COM/DCOM, Java™ RMI, or CORBA.

Road Map and Schedule

The Storage Networking Industry Association’s (SNIA) Storage Network Management Working Group (SNMWG) is chartered to identify, define, and

support open standards needed to address the increased management requirements imposed by storage area network environments. The Fibre Channel sub-group of the SNMWG is further focused to address the management of Fibre Channel Storage Area Networks. During the next year, the SNMWG-FC working group intends to provide the following:

1. Coordinate efforts with the other sub-groups of the SNMWG;
2. Coordinate the development of appropriate Fibre Channel Management Information Bases (MIBs) and map them into CIM;
3. Develop a SAN Management Architecture specification;
4. Provide implementation guides based on the SAN Management Architecture specification;
5. Standardize the host bus adapter (HBA) management in CIM.

Supplemental Reading

Basic Systems Management Disciplines

In simple terms, systems management is a set of disciplines that allows resources to be monitored, controlled, and serviced. Examples of resources in a storage area network include servers, disk drives, adapters, bridges, hubs, switches, storage controllers, and enclosures.

The basic systems management disciplines are:

- Event and fault management;
- Configuration management;
- Change management;
- Asset and accounting management;
- Performance management;
- Security management.

Fault management is concerned with the detection and processing of exceptional conditions associated with a resource. Event management deals with notifications of state changes in general, of which faults are one type. Both event and fault management includes the functions of notification, logging, and reporting. Faults additionally require problem isolation and determination, the possible execution of diagnostics, and problem resolution.

Configuration management is concerned with the monitoring and modification of a resource's setup or operational state. Configuration management also includes auto-discovery and directory mechanisms to support dynamic environments.

Change management is concerned with adding, updating, and deleting physical and logical components of a resource. Change management also provides the processes by which hardware, software, and firmware changes are planned and controlled.

Accounting and asset management are concerned with the identification and location of resources, measuring the use of a resource, and applying the appropriate charges to the consumer for the use of that resource.

Performance management provides the functions that allow the measurement of resources to determine their availability, utilization, and overall behavior. Performance management functions also provide for measurement data to be analyzed and reported.

Security management is concerned with managing the security infrastructure associated with the managed resource.

Systems Management Operations

Systems management operations may be broken down into three main categories:

- **Monitor Operations** - Monitor operations allow the state of managed components and their relationships to be determined. One typical example of monitoring functionality is querying the availability of a path to a storage device, or reading SAN port statistics.
- **Control Operations** - Control operations provide the ability to modify the state of a managed resource. Examples of this type of operation would be configuring a zone within a Fibre Channel network, or configuring Fibre Channel port characteristics.
- **Service Operations** - Service operations allow changes to be made to managed resources to fix faults or for preventative measures. For example, when SAN performance degrades due to a node or component fault, SAN management tools must be able to identify, isolate and correct the fault.

Policy Management

Policies represent the translation of business goals and objectives to rules and actions. The latter involve the definition of the rules and conditions which, when true, cause the policy actions to be executed. Actions can either result in some activity being performed or the cessation of activity. Alternately, actions can authorize or deter access.

To achieve desired business objectives, management operations are effected according to the criteria and rules defined by policy managers. For example, in order to meet service level agreements, service providers could rely on underlying functionality in the managed components to perform auto-configuration, self-diagnosis and self-correction of faults. This functionality could be invoked upon detection of the violation of policy conditions. Policy based management enables a proactive management environment versus the traditional reactive approaches to systems management.

Additional Information

Additional information regarding Fibre Channel SAN Management and CIM can be obtained at the following WEB sites:

<http://www.snia.org/>

<http://www.sansolutions.com/>

<http://www.dmtf.org/>