

UNCLASSIFIED



STORAGE AREA NETWORK SECURITY TECHNICAL IMPLEMENTATION GUIDE (STIG) OVERVIEW

Version 2, Release 5

30 January 2025

Developed by DISA for the DOD

UNCLASSIFIED

Trademark Information

Names, products, and services referenced within this document may be the trade names, trademarks, or service marks of their respective owners. References to commercial vendors and their products or services are provided strictly as a convenience to our users, and do not constitute or imply endorsement by the Defense Information Systems Agency (DISA) of any nonfederal entity, event, product, service, or enterprise.

TABLE OF CONTENTS

	Page
1. INTRODUCTION.....	1
1.1 Executive Summary.....	1
1.2 Authority.....	1
1.3 Vulnerability Severity Category Code Definitions.....	1
1.4 STIG Distribution.....	2
1.5 Document Revisions.....	2
1.6 Other Considerations.....	2
1.7 Product Approval Disclaimer.....	3
2. CONCEPTS AND TERMINOLOGY CONVENTIONS.....	4
2.1 Introduction	4
2.2 Architecture.....	4
2.3 SAN Security Concepts	6
2.3.1 Zoning.....	6
2.3.2 LUN Masking.....	8
2.4 Network and Host Security	8
2.4.1 Securing the Fabric Switch-to-Switch Connection.....	8
2.4.2 Securing the Management Interface.....	9
2.4.3 Securing the Host-to-Fabric Connection	9
2.5 Data Backup and Disaster Recovery	9

LIST OF TABLES

	Page
Table 1-1: Vulnerability Severity Category Code Definitions	2

LIST OF FIGURES

	Page
Figure 2-1. Sample SAN Architecture	5

1. INTRODUCTION

1.1 Executive Summary

The Storage Area Network (SAN) Overview is published as part of the Sharing Peripherals across the Network (SPAN) Security Technical Implementation Guide (STIG). The Storage Area Network (SAN) Checklist provides the technical security policies, requirements, and implementation details for applying security concepts to SAN technology. The SAN checklist is meant for use in conjunction with the Network Infrastructure and appropriate operating system (OS) STIGs.

This overview document supersedes the Sharing Peripherals across the Network Overview, Version 2 Release 2, 20 October 2012.

1.2 Authority

Department of Defense Instruction (DODI) 8500.01 requires that “all IT [information technology] that receives, processes, stores, displays, or transmits DOD information will be [...] configured [...] consistent with applicable DOD cybersecurity policies, standards, and architectures.” The instruction tasks that DISA “develops and maintains control correlation identifiers (CCIs), security requirements guides (SRGs), security technical implementation guides (STIGs), and mobile code risk categories and usage guides that implement and are consistent with DOD cybersecurity policies, standards, architectures, security controls, and validation procedures, with the support of the NSA/CSS [National Security Agency/Central Security Service], using input from stakeholders, and using automation whenever possible.” This document is provided under the authority of DODI 8500.01.

Although the use of the principles and guidelines in these SRGs/STIGs provides an environment that contributes to the security requirements of DOD systems, applicable NIST SP 800-53 cybersecurity controls must be applied to all systems and architectures based on the Committee on National Security Systems (CNSS) Instruction (CNSSI) 1253.

1.3 Vulnerability Severity Category Code Definitions

Severity Category Codes (referred to as CAT) are a measure of vulnerabilities used to assess a facility or system security posture. Each security policy specified in this document is assigned a Severity Category Code of CAT I, II, or III.

Table 1-1: Vulnerability Severity Category Code Definitions

Category	DISA Category Code Guidelines
CAT I	Any vulnerability, the exploitation of which will directly and immediately result in loss of Confidentiality, Availability, or Integrity.
CAT II	Any vulnerability, the exploitation of which has a potential to result in loss of Confidentiality, Availability, or Integrity.
CAT III	Any vulnerability, the existence of which degrades measures to protect against loss of Confidentiality, Availability, or Integrity.

1.4 STIG Distribution

Parties within the DOD and federal government's computing environments can obtain the applicable STIG from the DOD Cyber Exchange website at <https://cyber.mil/>. This site contains the latest copies of STIGs, SRGs, and other related security information. Those without a Common Access Card (CAC) that has DOD Certificates can obtain the STIG from <https://public.cyber.mil/>.

1.5 Document Revisions

Comments or proposed revisions to this document should be sent via email to the following address: disa.stig_spt@mail.mil. DISA will coordinate all change requests with the relevant DOD organizations before inclusion in this document. Approved changes will be made in accordance with the DISA maintenance release schedule.

1.6 Other Considerations

DISA accepts no liability for the consequences of applying specific configuration settings made on the basis of the SRGs/STIGs. It must be noted that the configuration settings specified should be evaluated in a local, representative test environment before implementation in a production environment, especially within large user populations. The extensive variety of environments makes it impossible to test these configuration settings for all potential software configurations.

For some production environments, failure to test before implementation may lead to a loss of required functionality. Evaluating the risks and benefits to a system's particular circumstances and requirements is the system owner's responsibility. The evaluated risks resulting from not applying specified configuration settings must be approved by the responsible AO. Furthermore, DISA implies no warranty that the application of all specified configurations will make a system 100 percent secure.

Security guidance is provided for the DOD. While other agencies and organizations are free to use it, care must be given to ensure that all applicable security guidance is applied at both the device hardening level and the architectural level due to the fact that some settings may not be configurable in environments outside the DOD architecture.

1.7 Product Approval Disclaimer

The existence of a STIG does not equate to DOD approval for the procurement or use of a product.

STIGs provide configurable operational security guidance for products being used by the DOD. STIGs, along with vendor confidential documentation, also provide a basis for assessing compliance with cybersecurity controls/control enhancements, which supports system assessment and authorization (A&A) under the DOD Risk Management Framework (RMF). Department of Defense AOs may request available vendor confidential documentation for a product that has a STIG for product evaluation and RMF purposes from disa.stig_spt@mail.mil. This documentation is not published for general access to protect the vendor's proprietary information.

AOs have the purview to determine product use/approval in accordance with (IAW) DOD policy and through RMF risk acceptance. Inputs into acquisition or pre-acquisition product selection include such processes as:

- National Information Assurance Partnership (NIAP) evaluation for National Security Systems (NSS) (<https://www.niap-ccevs.org/>) IAW CNSSP #11.
- National Institute of Standards and Technology (NIST) Cryptographic Module Validation Program (CMVP) (<https://csrc.nist.gov/groups/STM/cmvp/>) IAW federal/DOD mandated standards.
- DODIN Approved Products List (APL) (<https://aplits.disa.mil/processAPList.action>) IAW DODI 8100.04.

2. CONCEPTS AND TERMINOLOGY CONVENTIONS

2.1 Introduction

According to the Storage Networking Industry Association (SNIA), a storage area network (SAN) is any high-performance network whose primary purpose is to enable storage devices to communicate with computer systems and with each other. This definition applies regardless of the interconnect technology used (Fibre Channel, Ethernet, or other). However, since most SAN implementations use Fibre Channel switch fabrics for interconnectivity, this initial version of the SPAN STIG will focus on securing Fibre Channel SANs.

Storage networks use various storage devices, such as disk arrays, tape drives, robotic libraries, and file servers. Although very costly, a SAN increases the availability, improves local area network (LAN) bandwidth usage, and increases accessibility of DOD data by linking multiple storage devices on a dedicated storage network and making the storage space available to distributed application servers and clients. However, as is often the case, the requirement of the user to access data quickly can conflict with the need to keep DOD data secure.

SANs are becoming a viable and even preferred solution for data management on the networks. SANs are an excellent way of centralizing data to provide high performing and easy to manage data access. The architectural design and configuration of a SAN must ensure that data is highly available and accessible. However, data and communications security must also be considered essential to SAN equipment selection, implementation, and management. The storage industry is witnessing a rapid increase in servers and storage considerations within SANs. Greater storage accessibility means that more access points to the data will exist. This includes LAN, campus network, Metropolitan Area Network (MAN), Wide Area Network (WAN), and wireless access to the stored data. More access points mean that more attention must be paid to protecting information from unauthorized access.

Attaching a network to a set of storage resources presents security risks, which were not present when storage was simply connected to a server. A key component to protecting the SAN is to enable the highest security settings available in the server. However, with a SAN, storage can be directly attached to, and extended over a public network, such as the Internet, thus circumventing traditional operating system settings. It is critical to ensure the integrity and confidentiality of the data by other means. Network security policies must consider protection of data while in storage and during transmission.

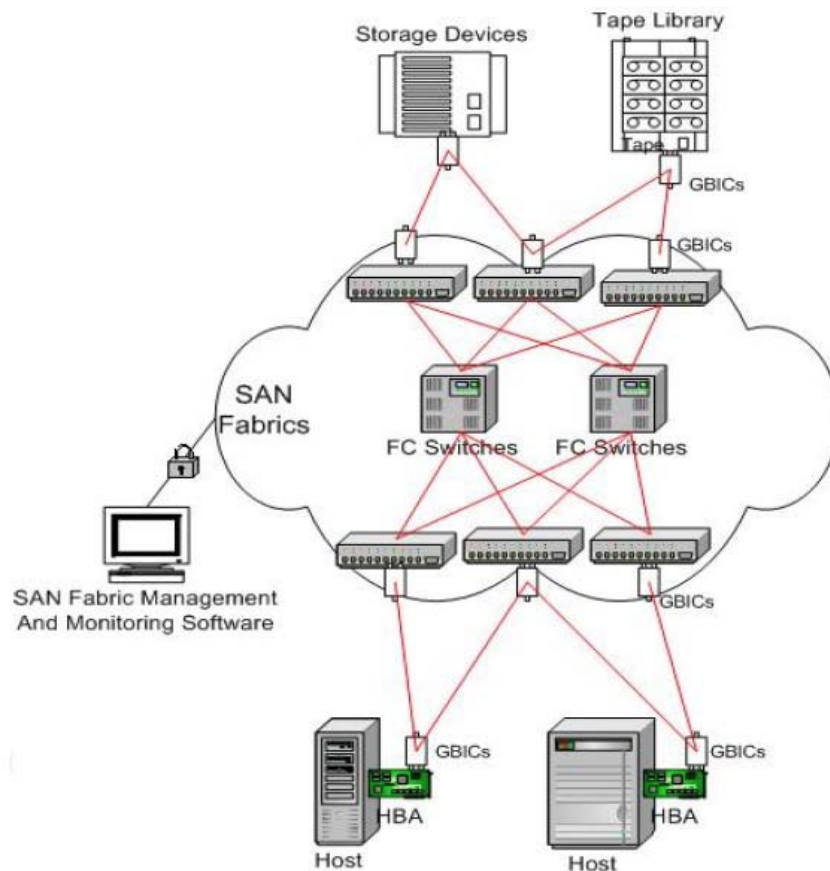
2.2 Architecture

This section discusses the typical components of the SAN architecture. SANs vary in complexity. Smaller SANs may not use all the components listed, while larger scaled installations may use a complex network for interconnecting multiple storage devices and server systems. Components may have both physical (port connections) and logical (zones) relationships to one another. See Figure 2-1, Sample SAN Architecture, obtained from the article, *Is your Storage Area Network Secure?* (Mohammed Haron, SANS Institute 2002).

The network connected storage environment may include some or all of the following components.

- SAN fabric (usually Fibre Channel switches).
- SAN fabric configuration management and monitoring software.
- SAN fabric security and access control software.
- Storage devices (disk arrays, tape devices, etc.).
- Hosts (application servers and client-level hosts).
- Host bus adapters (HBAs) which attach the hosts to the fabric switch.
- Network interface cards (NICs).
- Connection cabling, connectors, and optical to electrical signal converters or Gigabit Interface Converter (GBICs).
- Remote mirroring and replication storage applications.
- Backup and recovery software.

Figure 2-1. Sample SAN Architecture



Since the Fibre Channel switches or SAN fabric form the intelligent foundation of the SAN, it is vital that the switches are chosen with security of the DOD environment in mind. The following areas will impact the sites ability to secure the SAN and thus, should be carefully considered prior to purchasing a SAN fabric.

- Choose switches that support open industry standards, whenever possible.
- Switches must be configurable for use in a fully redundant architecture.
- Management and monitoring software should be robust, easy to use, and should support security protocols for secure configuration from a management server.
- Switches must have access control capabilities, such as passwords and configurable guards against configuration changes.
- Switches must have an Application Program Interface (API) so security applications from authorized DOD sources may be integrated into the SAN.

2.3 SAN Security Concepts

Fibre Channel continues to grow as the architecture of choice for providing high-speed, robust, and scalable interconnects for SANs. Fibre Channel enables the separation of storage and server, unlike the small computer system interface (SCSI), where the interconnect scheme is confined to the servers' cabinetry. A host of new security challenges consists of the exposure of critical business data to increased distances, greater availability, heterogeneous implementations, automatic re-configuration, increased services, and changes in strong model administration.

Without security, clients and application servers could see and use all devices attached to the SAN across the network connection making the devices more vulnerable to an attack. The storage administrator must make sure that users are only accessing and aware of the files to which they are authorized access. The two most common methods of providing access control security at this level in a SAN environment are zoning and Logical Unit Number (LUN) masking.

2.3.1 Zoning

Zoning separates the SAN into sub networks. The storage administrator uses zoning to group each host and storage device, then associates security and access policies to each group, thus preventing groups of devices from seeing or interacting with each other. Zoning may be used to group devices according to operating system, application, function, physical location, or other criteria as needed. SAN architectures provide port-to-port connections among servers and storage devices through bridges, switches, and hubs. Zoning is an efficient method of managing, partitioning, and controlling pathways to and from storage devices on the SAN fabric; as a result, storage resources are maximized, and data integrity and data security are maintained.

Note: For the Brocade switch, once zoning is activated, any device not explicitly placed in a zone is isolated and cannot communicate with other devices.

2.3.1.1 Hard Zoning

Hard zoning is accomplished by linking ports on the fabric through use of hardware and software. The specific implementation of hardware-enforced zoning differs depending on the switch vendor. Also, some restrictions may apply if the SAN environment consists of multiple vendor products (mixed-mode). For a specific zone, hard zoning can use Port World Wide Names (PWWN) and/or

port numbers to specify each device. Fabric switches and host bus adapters store and maintain copies of the zone's access control lists (ACLs) to verify access and routing prior to data transfers. Hard zoning requires each device pass through the switch's route table. For example, if two ports are not authorized to communicate with each other, the address will not appear in the route table and communication between those ports is blocked. The devices are physically unable to communicate with devices in another hard zone.

The exact details of how hard zoning is accomplished can differ significantly by vendor. Multi-vendor environments containing switches from different vendors will also impact zoning implementation. Rigorous testing must be employed to ensure that zoning works as intended, with repeated testing of all zones after adding or deleting switches.

The zoning process uses the ACL in the switch as the primary source of port numbers and worldwide names for each hard zone. The ACL is updated and propagated to other switches within the zone when changes are made to the zone. However, the HBA on the initiating devices also store a copy of the ACL. It is possible for the zoning information stored on the HBA to include old addresses, which are no longer allowed in the newly established zoning rules. The HBA's memory is non-persistent, thus a good practice is to force a state change update in the affected HBAs immediately after making zoning changes.

Zoning by ports is easier to implement, but less flexible than zoning by World Wide Name (WWN). Hard zoning does not allow zones to overlap or "follow" a zone member (device) that has its switch port changed. In other words, the zones need to be reconfigured whenever a Fibre Channel device in the SAN changes its switch port when hard zoning is used. When soft zoning is moved from one port to another, soft zoning remains associated with the device.

2.3.1.2 Soft Zoning

Zoning can also be implemented using Simple Name Server (SNS) software that runs on the fabric switch. By using the Node World Wide Name (NWWN) and/or the PWWN, soft zoning allows members of the zone to be defined. When a host logs into the SAN and requests available storage devices, the SNS will check the zoning table for all storage devices available for that host and the host will only see those devices that have been defined in the zoning table. However, it may be possible with certain operating systems for an attacker to bypass the SNS and go directly to the storage device thus soft zoning may present a potential security issue.

2.3.1.3 Configuring Zoning Components

Zone configurations are based on either the WWN of the device or the physical port that the device plugs into. Zoning components include zones, zone members, and zone sets.

A zone is made up of servers and storage arrays on the SAN fabric that can access each other through managed port-to-port connections. Devices in the same zone can recognize and communicate with each other, but not necessarily with devices in other zones, unless a device in that zone is configured for multiple zones.

Zone members are devices within the same assigned zone. Zone member devices are restricted to intra-zone communications, meaning that these devices can only interact with members within their assigned zone. Unless a device is configured for multiple zones, a zone member interacting with devices outside its assigned zone is not permitted.

Port number or WWN recognizes zone members. A WWN is a 64-bit number that uniquely identifies zone members.

Note: Information on soft zoning is provided to educate the reader on potential system vulnerabilities. Soft zoning is not recommended for use to protect access to SANs used to store sensitive DOD data.

2.3.2 LUN Masking

Many administrators use LUN masking to limit access to storage devices to further protect the SAN. LUN masking is a method of masking the unit number associated with the disk array. It is configured at the server console using the masking utility provided with the HBA driver. A single large disk array device can be sub-divided to serve a number of different hosts that are attached to the disk array through the SAN fabric. However, unlike zoning, the storage administrator can also limit access to each individual LUN inside the disk array by assigning specific LUNs to specific server(s).

LUN masking can be done either behind the disk array port or at the server HBA. It is more secure to mask LUNs at the disk array device, but not all disk array devices have LUN masking capability.

By filtering access to certain storage resources on the SAN, LUN masking goes one step beyond zoning alone. Also, by using a piece of code residing on each host connected to the SAN, LUN masking can be provided through hardware (i.e., intelligent bridges, routers, or storage controllers) or software. LUN masking effectively masks off the LUNs that are not assigned to the application server, allowing only the assigned LUNs to appear to the application server's operating system.

Managing paths by LUN masking is a reasonable solution for smaller SANs; however, this method requires an extensive amount of configuration and maintenance and is not recommended for larger SANs with large number of hosts or LUNs on the storage array. The complexity of maintaining this method of access control may present a security issue, as it is unlikely that the storage administrator will maintain the configuration for a large SAN.

2.4 Network and Host Security

This section contains general security policies, which apply to SAN network devices in the enclave and represent general best practices in any data-networking environment.

2.4.1 Securing the Fabric Switch-to-Switch Connection

This section discusses policies for securing the interconnection between fabric switches. A switch may attempt to illegally join a fabric or change the fabric topology. This is usually accomplished by having physical access to the SAN fabric. However, a management interface may enable this as well

from a remote location. An unauthenticated switch may be able to change the layout of the environment or cause denial of resource access to legitimate users.

2.4.2 Securing the Management Interface

To ensure that a trusted and secure management console-to-fabric communications layer exists, management-to-fabric technologies can use PKI and other encryption technologies. Not all SANs will have a dedicated management console, but the following checks should be applied to any host used to connect to the fabric for the purpose of managing the fabric devices such as the switches. PKI and other encryption help ensure that the management console or framework used to control the fabric is authentic and authorized. In addition, encryption methodologies can restrict the number of switches on the fabric from which management and configuration changes are propagated to the rest of the fabric.

2.4.3 Securing the Host-to-Fabric Connection

This section addresses security policies and technologies associated with the connection between the host servers (via the associated HBA) and the fabric switches. The goal is to secure this type of connection by explicitly allowing only authorized Fibre Channel HBAs of authorized hosts. All other HBAs are not allowed to attach to the port by default. Zoning to protect the host-to-fabric security technologies uses ACLs in much the same way that routers use ACLs. Thus, security concepts, which apply to router ACLs, must also be applied in securing the SAN environment. Enforcement of access control on each port by using zoning prevents unauthorized and intruder hosts from attaching to the fabric via any port. These restrictions may also be based on the source and destination addresses of the Internet Protocol (IP) packet as well as the service type (e.g., Simple Mail Transfer Protocol [SMTP], email, and Hypertext Transfer Protocol [HTTP]).

Typically, network facilities based on traditional networks provide connectivity between end-user platforms and server system components. It is also possible to connect end-user platforms directly to the Fibre Channel network, allowing the client host to directly access storage devices.

2.5 Data Backup and Disaster Recovery

Data backups in most network environments can be a major issue as the backup process may adversely impact network performance. Frequent backups during the day are often not an option for this reason. With a SAN, backup operations take place independent of the local network, making real-time, high-speed backup viable. Backup and recovery procedures are critical to the security and availability of the SAN system. If a system is compromised, shut down, or otherwise not available for service, this could hinder the availability of resources to the users.

After disasters, traditionally data recovery is handled using data tapes to restore the data. However, the use of SANs allows for various methods of automated data backup configurations which increase the availability of data in the event of a loss of the primary data image. Efficiently maintaining a redundant data image requires a low latency, high availability network infrastructure for which today's storage networks are suitable. When access to the primary data image (array or tape library) is lost to the primary site, the SAN fabric can be configured to automatically fail over to

the backup image. This backup image may be maintained on different hardware or on the same physical hardware, but a different logical drive.

The SAN switches will update their internal path information to provide alternate connectivity to storage devices. End user applications are not aware of this alternate routing. Failover in the event of interconnect or controller failure is essential when planning to consolidate an organization's data into a SAN because this method provides data availability, one of the essential elements of information assurance. Configuration of this feature is vendor specific, may require expensive hardware and management software, and is not available from all hardware vendors.