

UNCLASSIFIED



FIREWALL SECURITY REQUIREMENTS GUIDE (SRG) TECHNOLOGY OVERVIEW

Version 2, Release 3

27 October 2022

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 non-Federal entity, event, product, service, or enterprise.

TABLE OF CONTENTS

	Page
1. INTRODUCTION.....	1
1.1 Executive Summary	1
1.1.1 Security Requirements Guides (SRGs)	1
1.1.2 SRG Naming Standards	2
1.2 Authority	3
1.2.1 Relationship to STIGs.....	3
1.3 Vulnerability Severity Category Code Definitions	3
1.4 SRG and STIG Distribution	4
1.5 Document Revisions	4
1.6 Other Considerations.....	4
1.7 Product Approval Disclaimer.....	5
2. ASSESSMENT CONSIDERATIONS.....	6
2.1 NIST SP 800-53 Requirements	6
2.2 General Procedures	6
3. CONCEPTS AND TERMINOLOGY CONVENTIONS	7
3.1 Firewall Guidance	7
3.2 Types of Firewalls.....	7
3.2.1 Packet Filtering Firewalls	8
3.2.2 Stateful Inspection Firewalls	8
3.3 Network Architecture.....	8
3.4 Firewall Security Policy and Rules	9
3.5 Account Services.....	9
3.6 Device Management Audit Logs vs. Firewall Application Logs.....	10

LIST OF TABLES

	Page
Table 1-1: Vulnerability Severity Category Code Definitions	4

1. INTRODUCTION

1.1 Executive Summary

This Firewall Security Requirements Guide (SRG) provides the technical security policies and requirements for applying security concepts to firewall implementations. This SRG is applicable to the network-level firewall that provides packet and stateful filtering of network traffic up to and including Layer 4 of the OSI model. Layer 7 firewalls (such as Application Layer, Web, and Database) and host-based firewalls are not within the scope of this SRG.

The focus of the Firewall SRG is the configuration of the firewall application to protect network traffic. The firewalls are often included as part of other boundary control devices (e.g., VPNs and routers); therefore, this document will often be used in conjunction with other security guides to secure a network device. If an attacker can compromise the firewall, it can be rendered useless in detecting subsequent attacks against other hosts; thus, the device itself is often the target of attack. Therefore, the Network Device Management (NDM) SRG (appliance versions) or the applicable Operating System SRG (hosted version) is also an essential part of the security review.

The core functionality of a firewall is to provide initial traffic inspection and filtering at the internal and external security boundaries of the network, allowing or disallowing access based on an examination of the traffic content. The firewall enforces the security and access control policy between networks by examining all inbound and outbound packets for security threats.

The variety of implementation architectures of the network firewall presents a challenge for assessing the scope and applicability of security controls. Firewalls can be implemented as hardware-based or software-based systems. Hardware firewalls are dedicated network appliances with preinstalled security software. Software firewalls, on the other hand, can usually be installed on any available server that is equipped with a general-purpose network operating system, such as Windows or Linux. Firewall functionality is also usually included in routers, switches, or other network devices. Firewalls are often combined with other functionalities, such as VPN, including IPsec VPN, Domain Name System (DNS), antivirus, routing, Intrusion Detection System (IDS), and Network Address Translation (NAT). The Firewall SRG addresses the firewall functionality only. If other functions are implemented by the firewall device, the security requirements for that functionality must also be applied using additional SRGs and/or Security Technical Implementation Guides (STIGs). SRG requirements must be considered depending on the specific implementation of the firewall and the location of the installation (e.g., regional enclave, enterprise data center, or local network).

1.1.1 Security Requirements Guides (SRGs)

SRGs are collections of requirements applicable to a given technology family. SRGs represent an intermediate step between Control Correlation Identifiers (CCIs) and STIGs. CCIs represent discrete, measurable, and actionable items sourced from Information Assurance (IA) controls defined in a policy, such as the National Institute of Standards and Technology (NIST) Special

Publication (SP) 800-53. STIGs provide product-specific information for validating and attaining compliance with requirements defined in the SRG for that product's technology area.

There are four core SRGs: Application, Network, Operating System, and Policy. Each addresses the applicable CCIs in the context of the technology family. Subordinate to the core SRGs, there are Technology SRGs developed to address the technologies at a more granular level.

This Firewall SRG is based on the Network SRG. This Firewall SRG contains general check and fix information that can be used for products for which STIGs do not exist.

The STIGs based on this SRG will provide the product-specific technical implementation guidance for that product. The STIG will contain the specific check and fix information for the product it covers.

SRG Hierarchy example:

Application SRG
/__Database SRG
/__MS SQL Server 2005 STIG

The SRG relationship and structure provides the ability to identify requirements that may be considered not applicable for a given technology family and provide appropriate justification. It also provides the structure to identify variations in specific values based on the technology family. These variations will be captured once and will propagate down to the Technology SRGs and then to the STIGs. This will eliminate the need for each product-specific STIG to address items that are not applicable.

1.1.2 SRG Naming Standards

In an effort to establish consistency across the SRGs, a naming standard for the Group Title and STIGIDs has been established.

Technology SRG Naming Standards

For Technology SRG Group Title and STIGIDs the following applies:

{Core SRG value}-{Technology SRG}-{5- or 6-digit numeric sequence number}

Examples:

SRG-NET-000001-RTR-000001
SRG-APP-000001-COL-000001
SRG-NET-000001-VVSM-00001
SRG-OS-000001-UNIX-000001

Checks/fixes will be included at this level in a general form. These checks and fixes will apply for any STIGs that are created for products that do not have product-specific check and fix guidance.

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.

1.2.1 Relationship to STIGs

The SRG defines the requirements for various technology families, and the STIGs are the technical implementation guidelines for specific products. A single SRG/STIG is not all-inclusive for a given system, which may include but is not limited to Database, Web Server, and Domain Name System (DNS) SRGs/STIGs. For a given system, compliance with all (multiple) SRGs/STIGs applicable to a system is required.

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

	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 SRG and 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
- DoD Unified Capabilities (UC) Approved Products List (APL) (<https://www.disa.mil/network-services/ucco>) IAW DoDI 8100.04

2. ASSESSMENT CONSIDERATIONS

2.1 NIST SP 800-53 Requirements

All applicable baseline technical NIST SP 800-53 requirements and security best practice requirements are included in this SRG.

CNSSI 1253 defines the required controls for DoD systems, based on confidentiality, integrity, and availability (baseline) of the given information system. In all cases, CNSSI 1253, along with required baselines, will serve as the policy requirement for any given asset or information system.

2.2 General Procedures

This SRG has procedures that are intended to provide appropriate evaluation and remediation functions for a typically configured system. These procedures are not product specific and are intended for use when a product-specific STIG is not available.

The Firewall SRG addresses the firewall functionality only. If other functions (e.g., router, VPN) are integrated into the same device, the security requirements for that functionality must also be applied using additional SRGs and/or STIGs.

Security review must also include the NDM SRG (for appliance-based firewalls) or the applicable Operating System SRG (for hosted firewalls).

Layer 7 proxy filters, also known as application layer gateways, application firewalls, database firewalls, or web application firewalls, are out of scope. Use the Application Layer Gateway (ALG) SRG or product-specific STIG for assessment of these products.

IPsec is a Layer 3 protocol. It is out of scope for this SRG to avoid redundancy with the IPsec VPN STIG.

3. CONCEPTS AND TERMINOLOGY CONVENTIONS

3.1 Firewall Guidance

Firewall implementation must only allow incoming communications from authorized sources routed to authorized destinations. The firewall implementation must be configured in accordance with the security configuration settings based on DoD security configuration or implementation guidance, including NSA configuration guides, Communications Tasking Orders (CTOs), and Directive-Type Memorandums (DTMs). Firewall rules must comply with those specified or implied in the Firewall SRG, PPSM CAL, vulnerability assessments, and USCYBERCOM J3 and the Combatant Commands, Services, Agencies, and Field Activities (CC/S/A/FA) J6 information control policies.

Firewall configuration includes the binding of interfaces, security zones, and security policies. Traffic enters the firewall by way of interfaces. Security zones are configured for one or more interfaces with the same security requirements for filtering data packets. A security zone implements a security policy for one or multiple network segments. These security policies must be applied to inbound traffic as it crosses the network perimeter and as it crosses internal security domain boundaries. Security filters are also needed for unusual/unauthorized activities or conditions, which may include unusual protocols, unusual protocol behavior, or ports and attempted communications from trusted zones to external addresses.

3.2 Types of Firewalls

Network-level firewalls can be divided into four basic types, depending on what mechanisms are used to pass traffic from one security zone to another. Each mechanism uses a different layer of the Open Systems Interconnect (OSI) model. Firewalls capable of inspecting only the lower layers are more limited in the types of forwarding criteria that may be configured. Generally, firewalls with forwarding mechanisms operating at the lower OSI layers are faster but provide less robust inspection capabilities. Generally, firewalls can be subdivided into the following categories:

- Packet filtering firewalls
- Stateful packet filtering firewalls
- Firewalls with application awareness (e.g., ALGs, Database, Web)
- Hybrid firewalls

The following sections discuss the packet and stateful firewalls since these are within the scope of this SRG. This SRG is applicable to the network-level firewall that provides packet and stateful filtering of network traffic up to and including Layer 4 of the OSI model. Layer 7 firewalls (such as Application Layer, Web, and Database) are not within the scope of this SRG.

Note: Security guidance for firewalls with application awareness (Layer 7) is located in the ALG SRG.

3.2.1 Packet Filtering Firewalls

A packet filtering firewall does not track the state of the traffic. It filters the traffic via a filter (i.e., rules set, access control list [ACL]). These firewalls can filter based on source address, destination address, and session attributes such as source and destination ports. For example, allowing only approved IP addresses through the network perimeter will control access to required ports and services.

Packet filtering firewalls allow a direct connection between the two endpoints. This type of packet screening is configured to allow or deny traffic between two networks without disrupting the client/server model. However, packet filtering alone does not provide robust protection requirements because of its limitations in examining upper-layer data and providing detailed log data. Also, packet filtering firewalls are an all-or-nothing approach. If ports are open, then they are open to all traffic passing through that port, which in effect leaves a security hole in the network.

Packet filtering firewalls are susceptible to three common exploits: IP spoofing, buffer overruns, and ICMP tunneling. In an IP spoofing attack, the attacker sends data with a fake source address that appears valid to the firewall. A buffer overrun typically occurs when data sizes inside a buffer exceed what was allotted. ICMP tunneling allows an attacker to insert data into a legitimate ICMP packet.

3.2.2 Stateful Inspection Firewalls

Stateful inspection firewalls maintain a state table that tracks the state of each connection. The state table is aware of the network and transport layer protocols and traffic attributes, so traffic can be filtered using security policies that leverage this information. When a TCP connection is established, a source port and a destination port pair become part of the session, allowing the source system to receive information from the destination system. The client source port should be some port number greater than 1023 and less than 16384. Stateful inspection firewalls solve the vulnerability of permitting all the high-numbered ports by creating a state table containing the outbound TCP connections and their associated high-numbered port.

The directory known as the state table is then used to validate inbound traffic. Entries are created only for TCP connections or UDP streams that satisfy a defined security policy. Stateful inspection examines only the headers of data packets, which contain information such as the senders' and receivers' addresses and the type of protocol and data contained in the packet payload. As a result, stateful inspection technology cannot tell the difference between valid and harmful data. Like packet filtering, stateful packet inspection does not break the client/server model and therefore allows a direct connection to be made between the two endpoints.

3.3 Network Architecture

Each network implementation must include a description of the architectural implementation of the applicable security requirements. Care must be taken to include all of the firewall and content-filtering capabilities since these devices must work together to prevent operational

issues. The firewall implementation may include firewall functionality, which is part of other network components. For example, most routers, VPNs, and switches contain firewall functionality. Organizations should use firewalls wherever their internal networks and systems interface with external networks and where security requirements vary among their internal networks.

A packet filtering firewall, such as a customer premise edge router, must be implemented to filter traffic from external networks such as the NIPRNet, SIPRNet, and Internet. This premise router is the first line of defense in a defense-in-depth firewall solution. The premise router can block certain attacks and filter ports, protocols, and services prior to filtering operations at higher layers of the OSI stack by other firewall technologies.

The downstream firewall provides stateful inspection and application-level security. A firewall can be placed in several locations to provide protection from attack. Each implementation will differ depending on several factors, including the sensitivity of the networks, the network infrastructure, and the type of network traffic. Firewalls are used primarily to protect the boundaries of a network, although at times they can be used to separate an internal security domain from the rest of an enclave (LAN to LAN).

3.4 Firewall Security Policy and Rules

The organization's firewall policy must be documented in the security plan. The policy specifies how the firewall implementation will handle network traffic based on specific characteristics of the traffic (e.g., IP address, IP address range, port, interface, protocol, session identification, application, and content type). Most firewalls implement this policy by blocking information flows based on a set of screening rules defined within the firewall devices. Rule sets are also known as Access Control Lists (ACLs). An action (e.g., deny, permit, or drop) is usually associated with each rule. Any traffic encountered that does not meet the criteria for a configured rule is denied by default.

The enclave firewall rules must be based on the applications being used within the enclave. All non-required ports and services will be blocked by the most restrictive rules possible. Traffic allowed through the firewall will be configured in accordance with applicable DoD policy. Perimeter filtering rules can be applied to any internal firewall device or router and should be implemented to the fullest extent possible to minimize internal threats. Firewall administrators must conduct periodic reviews of firewall rules, policies, and configuration to ensure that the firewall remains effective against new attack techniques or changes in the network environment. Removing unnecessary or outdated rules and configuration increases the security and operational efficiency of the firewall.

3.5 Account Services

Firewall applications do not use specific accounts other than for administrative purposes. This requirement is applicable for temporary accounts created or maintained using the firewall application itself rather than the underlying operating system or an authentication server. Accounts created and maintained on authentication servers or Authentication, Authorization, and

Accounting (AAA) devices (e.g., RADIUS, LDAP, or other authentication services) are secured using the applicable SRG or STIG.

Security for the operating system or authentication server accounts is beyond the scope of this SRG. Account management requirements in the SRG apply to accounts created and managed on or by the firewall components. These requirements are often managed by the underlying operating system and thus are not part of the Firewall SRG scope, nor are they intended to be part of the network monitoring requirements of the firewall implementation.

3.6 Device Management Audit Logs vs. Firewall Application Logs

Two types of log files are required for each component of the organization's firewall implementation: the device management (backplane) audit log and firewall application log. The audit log stores the results of enforcement actions based on the access control restrictions, use of user privileges, and other security policies. This type of functionality is usually performed by the OS on network devices. The application log stores events detected as part of the firewall content-filtering activity. Depending on the implementation, these logs may be separate or combined. However, a good security practice is to have the capability to separate the logs, thus allowing the logs to be aggregated into separate databases.

The best practice is for the organization to use a central logging server (e.g., SYSLOG) to aggregate audit logs. Therefore, many of the audit requirements are not part of the Firewall SRG scope. However, the firewall implementation must also have the capability to centralize the application log. Application log requirements are separate from audit log requirements within the Firewall SRG.