Juniper Networks Security Appliances
Security Target

Version 0.8
April 6, 2012

Prepared for:
Juniper Networks

1194 North Mathilda Ave
Sunnyvale, CA 94089-1206

Prepared By:
Science Applications International Corporation

Common Criteria Testing Laboratory

6841 Benjamin Franklin Drive
Columbia, MD 21046

LIST OF TABLES

1 Security Target Introduction

This section identifies the Security Target (ST) and Target of Evaluation (TOE) identification, ST conventions, ST conformance claims, and the ST organization. The TOE is any of the Security Appliances provided by Juniper Networks and identified in section 1.1. The Target of Evaluation (TOE) primarily supports the definition of and enforces information flow policies among network nodes. The TOE provides for stateful inspection of every packet that traverses the network. All network traffic from one network zone to another or between two networks within the same network zone passes through the TOE, if the network and appliances are properly connected and configured. Information flow is controlled on the basis of network node addresses, protocol, and services requested. In support of the information flow security functions, a security appliance ensures that security relevant activity is audited, that its own functions are protected from potential attacks, and provides the security tools to manage all of the TOE's security functions.

The Security Target contains the following additional sections:

- TOE Description (Section 2)
- Security (Section 3)
- Security Objectives (Section 4)
- IT Security Requirements (Section 5)
- TOE Summary Specification (Section 6)
- Protection Profile Claims (Section 7)
- Rationale (Section 8)

1.1 Security Target, TOE and CC Identification

ST Title – Juniper Networks Security Appliances Security Target

ST Version – Version 0.8

ST Date – April 6, 2012

TOE Identification – The TOE consists of one or more of the following security appliances running the specified ScreenOS firmware version:

Product	Part Numbers	Firmware Version
Juniper Networks NetScreen ISG 1000	NS-ISG-1000, NS-ISG-1000-DC, NS-ISG-1000B, NS-ISG-1000B-DC	6.3.0r6
Juniper Networks NetScreen ISG 2000	NS-ISG-2000, NS-ISG-2000-DC, NS-ISG-2000B, NS-ISG-2000B-DC	6.3.0r6
Juniper Networks NetScreen 5200	NS-5200, NS-5200-DC	6.3.0r6
Juniper Networks NetScreen 5400	NS-5400, NS-5400-DC	6.3.0r6
Juniper Networks SSG5 Secure Services Gateway	SSG-5-SB, SSG-5-SH	6.3.0r6

Juniper Networks SSG20 Secure Services Gateway	SSG-20-SB, SSG-20-SH	6.3.0r6
Juniper Networks SSG140 Secure Services Gateway	SSG-140-SB, SSG-140-SH	6.3.0r6
Juniper Networks SSG320M Secure Services Gateway	SSG-320M-SH, SSG-320M-SH-N-TAA, SSG-320M-SH-DC-N-TAA	6.3.0r6
Juniper Networks SSG350M Secure Services Gateway	SSG-350M-SH, SSG-350M-SH-N-TAA, SSG-350M-SH-DC-N-TAA	6.3.0r6
Juniper Networks SSG520M Secure Services Gateway	SSG-520M-SH, SSG-520M-SH-N-TAA, SSG-520M-SH-DC-N-TAA	6.3.0r6
Juniper Networks SSG550M Secure Services Gateway	SSG-550M-SH, SSG-550M-SH-N-TAA, SSG-550M-SH-DC-N-TAA	6.3.0r6

TOE Developer – Juniper Networks

Evaluation Sponsor – Juniper Networks

CC Identification – Common Criteria for Information Technology Security Evaluation, Version 3.1 Revision 3, July 2009

1.2 Conformance Claims

This ST and the TOE it describes are conformant to the following CC specifications:

- Common Criteria for Information Technology Security Evaluation Part 2: Security Functional Requirements, Version 3.1, Revision 3, July 2009
    - Part 2 Conformant
- Common Criteria for Information Technology Security Evaluation Part 3: Security Assurance Requirements, Version 3.1, Revision 3, July 2009
    - Part 3 Conformant
    - Assurance Level: EAL2 augmented with ALC_FLR.2

The TOE meets all of the security requirements of the following Protection Profiles:

- U.S. Government Protection Profile for Traffic-Filter Firewall in Basic Robustness Environments, Version 1.1, July 25, 2007 (TFFW BR).

1.3 Conventions

This security target is claiming compliance with a Protection Profile (PP).

The conventions used in this security target are intended to highlight the completion of operations made within this security target. While this security target will include the operations made by the PP upon the CC requirements it is not the author's intent to highlight those operations (i.e., use bold, italics or special fonts). Therefore, keywords (e.g., selection, assignment and refinement) and formatting (e.g., special fonts) that may be used within the PP to designate operations made by the PP are being removed within this ST. The brackets used by the PPs to designate operations completed by the PP are left in the requirements.

The following conventions have been applied to indicate operations that this ST is making to the requirements in the TFFW BR Protection Profile:

- Security Functional Requirements – Part 2 of the CC defines the approved set of operations that may be applied to functional requirements: iteration, assignment, selection, and refinement.
    - Iteration: allows a component to be used more than once with varying operations. In the ST, iteration is indicated by a number in parentheses placed at the end of the component. For example FDP_ACC.1(1) and FDP_ACC.1(2) indicate that the ST includes two iterations of the FDP_ACC.1 requirement, (1) and (2).
    - Assignment: allows the specification of an identified parameter. Assignments are indicated using bold and are surrounded by brackets (e.g., [**assignment**]). Note that an assignment within a selection would be identified in italics and with embedded bold brackets (e.g., [***[selected-assignment]***]).
    - Selection: allows the specification of one or more elements from a list. Selections are indicated using bold italics and are surrounded by brackets (e.g., [***selection***]).
    - Refinement: allows the addition of details. Refinements are indicated using bold, for additions, and strike-through, for deletions (e.g., "… **all** objects …" or "… ~~some~~ **big** things …").
- Other sections of the ST – Other sections of the ST use bolding to highlight text of special interest, such as captions.

1.4 Terminology

Address	The network portion of an IP address. Most IP addresses have a network portion and a node portion.
Address Shifting	A mechanism for creating a one-to-one mapping between any original address in one range of addresses and a specific translated address in a different range.
Application-Specific Integrated Circuit (ASIC)	A customized microchip, which is designed for a specific application.
Authorized Administrator	A role which human users may be associated with to administer the security parameters of the TOE. Such users are not subject to any access control requirements once authenticated to the TOE and are therefore trusted to not compromise the security policy enforced by the TOE.
Authorized external server	Any IT product or system, outside the scope of the TOE that may administer the security parameters of the TOE. Such servers are not subject to any access control requirements once authenticated to the TOE and are therefore trusted to not compromise the security policy enforced by the TOE.
Central Processing Unit (CPU)	The CPU controls the operation of a computer.
Destination Network Address Translation (NAT-dst)	The translation of the original destination IP address in a packet header to a different destination address. ScreenOS supports the translation of one or several original destination IP addresses to a single IP address ("one-to-one" or "many-to-one" relationships). The TOE also supports the translation of one range of IP addresses to another range (a "many-to-many" relationship) using address shifting. When the TOE performs NAT-dst without address shifting it can also map the destination port number to a different predetermined port number. When the TOE performs NAT-dst with address shifting, it cannot also perform port mapping.

Dynamic IP (DIP) Pool	A dynamic IP (DIP) pool is a range of IPv4 addresses from which the security appliance can dynamically or deterministically take addresses to use when performing network address translation on the source IPv4 address (NAT-src) in IP packet headers.
Dynamic Random Access Memory (DRAM)	A type of computer memory that is stored in capacitors on a chip. Most computers have DRAM chips, because they provide a lot of memory at a low cost.
External server	Any IT product or system, untrusted or trusted, outside of the TOE that interacts with the TOE.
Federal Information Processing Standards (FIPS)	The Federal Information Processing Standards Publication (FIPS PUB) series issued by the U.S. National Institute of Standards and Technology as technical guidelines for U.S. Government procurements of information processing system equipment and services.
FIPS 140-2	The U.S. Government standard for security requirements to be met by a cryptographic module used to protect unclassified information in computer and communication systems. The standard specifies four increasing levels (from 'Level 1' to 'Level 4') of requirements to cover a wide range of potential applications and environments. The requirements address basic design and documentation, cryptographic module ports and interfaces, authorized roles and services, physical security, operational environment, cryptographic key management, electromagnetic interference and electromagnetic compatibility (EMI/EMC), and self-testing.
Firmware	Software stored in Read Only Memory (ROM) or Programmable Read-Only Memory(PROM) essential programs that remain even when the system is turned off. Firmware is easier to change than hardware but more permanent than software stored on disk.
Flash Memory	A small printed circuit board that holds large amounts of data in memory. Flash memory is used because it is small and holds its data when the computer is turned off.
Hyper Text Transfer Protocol (HTTP)	The protocol most commonly used in the World-Wide Web to transfer information from Web servers to Web browsers.
Internet Control Message Protocol (ICMP)	An extension to the Internet Protocol (IP), which is used to communicate between a gateway and a source host, to manage errors and generate control messages.
Mapped IP Address (MIP)	A MIP is a direct one-to-one mapping of traffic destined from one IP address to another IP address. The TOE forwards incoming traffic destined for a MIP to the host with the address to which the MIP points. Essentially, a MIP is static destination address translation, mapping the destination IP address in an IP packet header to another static IP address. When a MIP host initiates outbound traffic, the TOE translates the source IP address of the host to that of the MIP address. This bidirectional translation symmetry differs from the behavior of source and destination address translation. MIPs allow inbound traffic to reach private addresses in a zone whose interface is in NAT mode. MIPs also provide part of the solution to the problem of overlapping address spaces at two sites connected by a VPN tunnel.

Network Address Translation (NAT)	NAT involves translating the source IP address in a packet header to a different IP address. In the case of a traditional NAT, the translated source IP address comes from the IP address of the egress interface. When the security appliance uses the IP address of the egress interface, it translates all original source IP addresses to the address of the egress interface.
NAT Source (NAT-src)	NAT-src involves translating the source IP address in a packet header to a different IPv4 address from a dynamic IPv4 (DIP) address pool. When the security appliance draws addresses from a DIP pool, it can do so dynamically or deterministically. When doing the former, it randomly draws an address from the DIP pool and translates the original source IPv4 address to the randomly selected address. When doing the latter, it uses address shifting to translate the source IPv4 address to a predetermined IPv4 address in the range of addresses that constitute the pool.
Network Basic Input/Output System (NetBIOS)	An application programming interface used in conjunction with other programs to transmit messages between applications running on PCs hooked to a local area network.
Network	A composition of a communications medium and components attached to that medium whose responsibility is the transfer of information. Such components may include automated information systems, packet switches, telecommunications controllers, distribution centers, technical management, and control devices. It is a set of devices such as computers, terminals, and printers that are physically connected by a transmission medium so that they can communicate with each other.
Node	A concentration point in a network where numerous trunks come together at the same switch.
Packet	A block of data sent over the network transmitting the identities of the sending and receiving stations, error-control information, and message.
Port Address Translation (PAT)	The translation of the original source port number in a packet to a different, randomly designated port number.
Public-Key Infrastructure (PKI)	A system of Certificate Authorities (CAs) (and, optionally, Registration Authorities (RAs) and other supporting servers and agents) that perform some set of certificate management, archive management, key management, and token management functions for a community of users in an application of asymmetric cryptography.
Session	A series of interactions between two communication end points that occur during the span of a single connection. Typically, one end point requests a connection with another specified end point and if that end point replies agreeing to the connection, the end points take turns exchanging commands and data ("talking to each other"). The session begins when the connection is established at both ends and terminates when the connection is ended.
Session Table	A resource within the security appliance that maintains a list of active sessions. The session table is utilized to verify if any requesting information flows may already have an established session.

Stateful inspection	Also referred to as *dynamic packet filtering*. Stateful inspection is a firewall mechanism that works at the network layer. Unlike static packet filtering, which examines a packet based on the information in its header, stateful inspection tracks each connection traversing all interfaces of the firewall and makes sure they are valid. An example of a stateful firewall may examine not just the header information but also the contents of the packet up through the application layer in order to determine more about the packet than just information about its source and destination. A stateful inspection firewall also monitors the state of the connection and compiles the information in a state table. Because of this, filtering decisions are based not only on administrator-defined rules (as in static packet filtering) but also on context that has been established by prior packets that have passed through the firewall. As an added security measure against port scanning, stateful inspection firewalls close off ports until connection to the specific port is requested.
Synchronous Dynamic Random Access Memory (SDRAM)	High-speed DRAM that adds a separate clock signal to the control signals. SDRAM can transfer bursts of non-contiguous data at 100 MBytes/sec, and has an access time of 8-12 nanoseconds. It comes in 64-bit modules: long 168-pin Dual In-line Memory Modules (DIMMs).
Tampering	An unauthorized modification that alters the proper functioning of equipment or a system in a manner that degrades the security or functionality it provides.
Transmission Control Protocol/Internetwork Protocol (TCP/IP)	A communications protocol developed under contract from the U.S. Department of Defense to internetwork dissimilar systems. Transport Control Protocol/Internet Protocol. Refers to the Internet Protocol Suite, which includes TCP and IP, as well as several other protocols, used by computers to communicate with each other. TCP/IP is the standard protocol used on the Internet. It can also be used as a communications protocol in the private networks called intranets and in extranets. TCP/IP is a two-layered protocol The higher layer, Transmission Control Protocol, manages the marshalling of a message or file into smaller packets that are transmitted over the Internet and received by a TCP layer that reassembles the packets into the original message. The lower layer, Internet Protocol, handles the address part of each packet so that it gets to the right destination.
Tunneling	Use of one data transfer method to carry data for another method.
User Datagram Protocol (UDP)	A communications protocol for the Internet network layer, transport layer, and session layer, which makes it possible to send a datagram message from one computer to an application running in another computer. Like TCP (Transmission Control Protocol), UDP is used with IP (the Internet Protocol). Unlike TCP, UDP is connectionless and does not guarantee reliable communication; the application itself must process any errors and check for reliable delivery.
Virtual IP Address (VIP)	A Virtual IP address (VIP) maps traffic received at one IP address to another address based on the destination port number in the packet header. In other words, the actual destination IP addresses for two VIPs can be the same, yet the TOE uses destination port number to determine where to forward traffic.

Virtual Private Network (VPN)	An Internet-based system for information communication and enterprise interaction. A VPN uses the Internet for network connections between people and information sites. It includes stringent security mechanisms so that sending private and confidential information is as secure as in a traditional closed system.
Virtual Router (VR)	A virtual router (VR) is the component of ScreenOS that performs routing functions. A virtual router functions as a router. It has its own interfaces and its own routing table. By default, a security appliance supports two virtual routers: Untrust-VR and Trust-VR. This allows the security appliance to maintain two separate routing tables and to conceal the routing information in one virtual router from the other. For example, the untrust-vr is typically used for communication with untrusted parties and does not contain any routing information for the protected zones. Routing information for the protected zones is maintained by the trust-vr. Thus, no internal network information can be gleaned by the surreptitious extraction of routes from the untrust-vr.
Virtual System	A virtual system (vsys) is a subdivision of the main system that appears to the user to be a stand-alone entity. Virtual systems reside separately from each other in the same security appliance. Each one can be managed by its own virtual system administrator. Virtual systems are outside the scope of the evaluated configuration of the TOE.
Zone(s)	A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or logical entity that performs a specific function (a function zone).

2 TOE Description

The Target of Evaluation (TOE) is a Juniper Networks Security Appliance as enumerated in the set of products listed in section 1.1.

Juniper Networks Security Appliances, hereafter referred to as security appliances, are network security devices designed and manufactured by Juniper Networks, 1194 North Mathilda Avenue, Sunnyvale, California 94089-1206 U.S.A, herein called simply Juniper. The security appliances are individual network security devices. Each security appliance can operate as a central security hub in a networked configuration. The security appliances control traffic flow between connected networks. The security appliances integrate stateful packet inspection firewall and traffic management features.

2.1 TOE Overview

Juniper's line of security appliances combines firewall, virtual private networking (VPN), and traffic management functions. Configuring and managing appliances is accomplished using a command line interface (CLI).

The TOE includes only the security appliance. Each security appliance runs ScreenOS, which is a custom operating system. The security appliance models listed in column 1 of the table in section 1.1 each consists of hardware and ScreenOS (which runs in firmware).

The security appliances use a technique known as 'stateful inspection' rather than an 'application proxy,' as stateful inspection offers the combination of security and performance. Stateful inspection firewalls examine each packet, and track application-layer information for each connection, by setting up a state table that spans multiple packets. This is used to determine whether incoming packets are legitimate. It eliminates the requirement to establish a TCP session with the firewall itself to access a service on the other side of the firewall (i.e. proxy the service).

To perform routing functions ScreenOS implements a virtual router (VR) component, which functions as a router and has its own interfaces and its own routing table. The TOE supports two predefined virtual routers, trust-vr and untrust-vr. This allows the security appliance to maintain two separate routing tables and to conceal the routing information in one virtual router from the other. For example, the untrust-vr is typically used for communication with untrusted parties and does not contain any routing information for the protected zones. Routing information for the protected zones is maintained by the trust-vr. Thus, no internal network information can be gleaned by the surreptitious extraction of routes from the untrust-vr. There are no limitations on the number of virtual routers that can be configured and used in the evaluated configuration.

2.2 Product Description

Juniper Networks Security Appliances all share a very similar hardware architecture and packet flow. All run ScreenOS with common core features across all products. All security appliances perform the same security functions and export the same types of interfaces. A sample of the differences between these products is listed below.

- The SSG 5 and SSG 20 use an Intel IXP625 ASIC; the SSG 140 uses the Intel IXP2325. The Intel IXP ASICs provide acceleration of AES, and SHA-1. The remaining cryptographic and firewall functionality is performed in software.
- The 320M, 350M, 520M and 550M use the Cavium Nitrox Lite ASIC to accelerate AES, SHA-1 and modular exponentiation operations. The remaining cryptographic and firewall functionality is performed in software.
- The Juniper Networks NetScreen-5200, NetScreen-5400, NetScreen-ISG1000 and NetScreen-ISG2000 use one or more custom GigaScreen3 ASICs. The GigaScreen3 ASIC is capable of providing most of the firewall and cryptographic functionality, and uses the CPU as a co-processor for handling management traffic and first packet inspections (policy lookups). The GigaScreen3 ASIC can process an incoming packet, perform a session lookup, NAT, TCP/IP sequence checking, and can then send the packet back out

of the device without ever being processed by the system CPU. The only time the CPU is used is for first packet inspection, management traffic, and packet fragment reassembly for inspection. These platforms use the Cavium Nitrox Lite ASIC for acceleration of modular exponentiation operations.

The remainder of this section will describe the security appliance hardware and ScreenOS. ScreenOS is described in terms of the policies enforced by the security appliance.

2.2.1 Hardware

The hardware is manufactured to Juniper's specifications by sub-contracted manufacturing facilities. Juniper's custom OS, ScreenOS, runs in firmware. The security appliances provide no extended permanent storage like disk drives and no abstractions like files. Audit information is stored in memory. The main components of a security appliance are the processor, ASIC, memory, interfaces, and surrounding chassis and components. The differences between security appliances are the types of processor(s), traffic interfaces, management interfaces, number of power supplies, type of ASIC, and redundancy to ensure high availability. The supported network interfaces that carry network traffic include support for Gigabit or 10/100Mbps copper-based connections as well as Fibre channel connections. All devices support 10/100Mbps ethernet connectivity, while some also provide a management interface through an RJ-45 serial port.

2.2.2 ScreenOS

ScreenOS powers the entire system. At its core is a custom-designed, real time operating system built from the outset to deliver security and performance. ScreenOS provides an integrated platform for its functions, including:

- Stateful inspection firewall
- Traffic management
- Site-to-Site VPN

ScreenOS does not support a general-purpose, computing environment.

2.2.2.1 Policies

Security appliances enforce information flow control decisions by defining policies which permit, deny, reject, NAT or tunnel information flows in accordance with the rules defined in each policy. These policies are applied by the TOE to all IPv4 traffic, IPv6[1] traffic was not included as part of the evaluation. All policies on a security appliance include the following attributes:

- Direction – The direction of traffic between two security zones (from a source zone to a destination zone);
- Source address – The address from which traffic initiates;
- Destination address – The address to which traffic is sent;
- Service – The type of traffic transmitted; and,
- Action – The action that the security appliance performs when it receives traffic meeting the first four criteria: permit, deny (drop silently), reject (drop with ICMP error), NAT (perform address translation), or tunnel (permit with encryption or decryption)[2].

Security appliances enforce policies based on a service. A service specifies the protocol (TCP or UDP), the port number, the service group, the timeout and the flag associated to a specific service and maps the service to a defined name.

Security appliances provide three different types of policies which support the information flow control decisions enforced by the TOE. This includes Interzone Policies, Intrazone Policies, and Global Policies. These policies are invoked when determining the appropriate decision to make on an information flow. The following sections describe differences between each of these three types of policies.

[1] The IPv6 implementation is unchanged between ScreenOS 6.2 and 6.3.
[2] Tunneling refers to the site-to-site VPN functionality described in this Security Target.

2.2.2.1.1 Interzone policies

Interzone policies provide traffic control between security zones. You can set interzone policies to permit, deny, or tunnel traffic from one zone to another. Using stateful inspection techniques, the TOE maintains a table of active TCP sessions and active UDP "pseudo" sessions so that it can allow replies to service requests.

2.2.2.1.2 Intrazone Policies

Intrazone policies provide traffic control between interfaces bound to the same security zone. The source and destination addresses are in the same security zone, but reached via different interfaces on the TOE. Like interzone policies, intrazone policies control traffic flowing unidirectionally. To allow traffic initiated at either end of a data path, you must create two policies—one policy for each direction.

Intrazone policies do not support VPN tunnels or source network address translation (NAT-src) when it is set at the interface level (set interface NAT). However, intrazone policies do support policy-based NAT-src and NAT-dst. They also support destination address translation when the policy references a mapped IP (MIP) as the destination address. A mapped IPv4 address is a direct one-to-one mapping of traffic destined for one IPv4 address to another IPv4 address.

2.2.2.1.3 Global Policies

Unlike interzone and intrazone policies, global policies do not reference specific source and destination zones. Global policies reference user-defined Global zone addresses or the predefined Global zone address "any". These addresses can span multiple security zones. For example, if you want to provide access to or from multiple zones, you can create a global policy with the Global zone address "any", which encompasses all addresses in all zones.

Because a TOE configured in transparent mode operates at layer-2 without a layer 3 IP address(see section 2.3.1), a VPN can only be established using zone based policies that do not require a layer 3 IP address (i.e., interzone policies). Thus, global policies are not enforced on Transparent Mode VPNs.

2.2.2.1.4 Order of Invocation

When the TOE initiates a policy lookup, it first checks to see if the security zones are the same or different. If the zones are different, the TOE performs a policy lookup in the interzone policy set list. If the zones match, the TOE performs a policy lookup in the intrazone policy set. If a policy is not found within either the interzone or intrazone set lists, the TOE performs a policy lookup in the global policy set list.

2.2.2.1.5 Firewall User Authentication

A firewall policy may require authentication prior to permitting traffic to cross the firewall. The authentication option may be combined with an interzone, intrazone, or global policy and requires a username and password to be provided via IKE, XAuth, L2TP, HTTP, FTP or Telnet. Successful authentication does not grant administrative access to the TOE.

2.3 TOE Configurations

The TOE provides three possible ways to configure a network interface. A network interface may be configured to operate in Transparent Mode, NAT Mode, or Route Mode. In addition, the TOE also supports Site-To-Site VPNs using a pre-shared key for authentication. These various configurations are further described below.

These interface modes each determine how packets are routed and filtered by the TOE. Each instance of the TOE can include one, a combination of, or all three interface modes. However, each individual network interface may only be configured with one interface mode and may not share a combination of or all three interface modes with one physical network interface. Each interface mode consistently satisfies all of the TOE security functional requirement claims identified in this ST. These three interface modes are further described below.

2.3.1 Transparent Mode

When a TOE interface is configured in Transparent Mode, the TOE filters packets traversing the firewall without modifying any of the source or destination information in the IP packet header. All interfaces behave as though they are part of the same network, with the TOE acting much like a Layer 2 switch or bridge. In Transparent mode, the IP addresses of interfaces are set at 0.0.0.0, making the presence of the TOE invisible, or "transparent," to users.

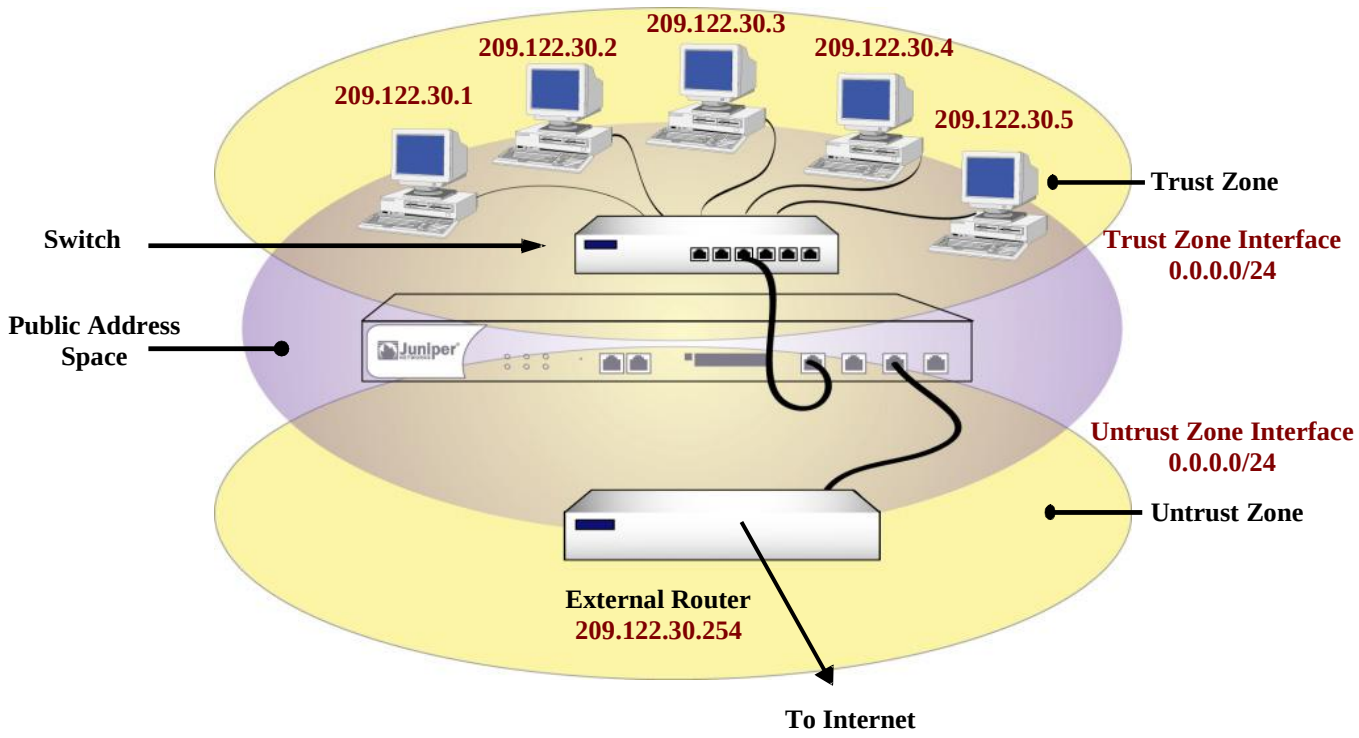

Figure 2.1: Transparent Mode

2.3.2 NAT Mode

When an ingress interface is in Network Address Translation (NAT) mode, the security appliance, acting like a Layer 3 switch (or router), translates two components in the header of an outgoing IP packet destined for the Untrust zone: its source IP address and source port number. The security appliance replaces the source IP address of the originating host with the IP address of the Untrust zone interface. Also, it replaces the source port number with another random port number generated by the security appliance.

When the reply packet arrives at the security appliance, the device translates two components in the IP header of the incoming packet: the destination address and port number, which are translated back to the original numbers.

The security appliance then forwards the packet to its destination. NAT adds a level of security not provided in Transparent mode: The addresses of hosts sending traffic through an ingress interface in NAT mode (such as a Trust zone interface) are never exposed to hosts in the egress zone (such as the Untrust zone) unless the two zones are in the same virtual routing domain and the security appliance is advertising routes to peers through a dynamic routing protocol (DRP). Even then, the Trust zone addresses are only reachable if you have a policy permitting inbound traffic to them. (If you want to keep the Trust zone addresses hidden while using a DRP, then put the Untrust zone in the untrust-vr and the Trust zone in the trust-vr, and do not export routes for internal addresses in the trust-vr to the untrust-vr.) If the security appliance uses static routing and just one virtual router, the internal addresses remain

hidden when traffic is outbound, due to interface-based NAT. The policies you configure control inbound traffic. If you use only mapped IP (MIP) and virtual IP (VIP) addresses as the destinations in your inbound policies, the internal addresses still remain hidden.

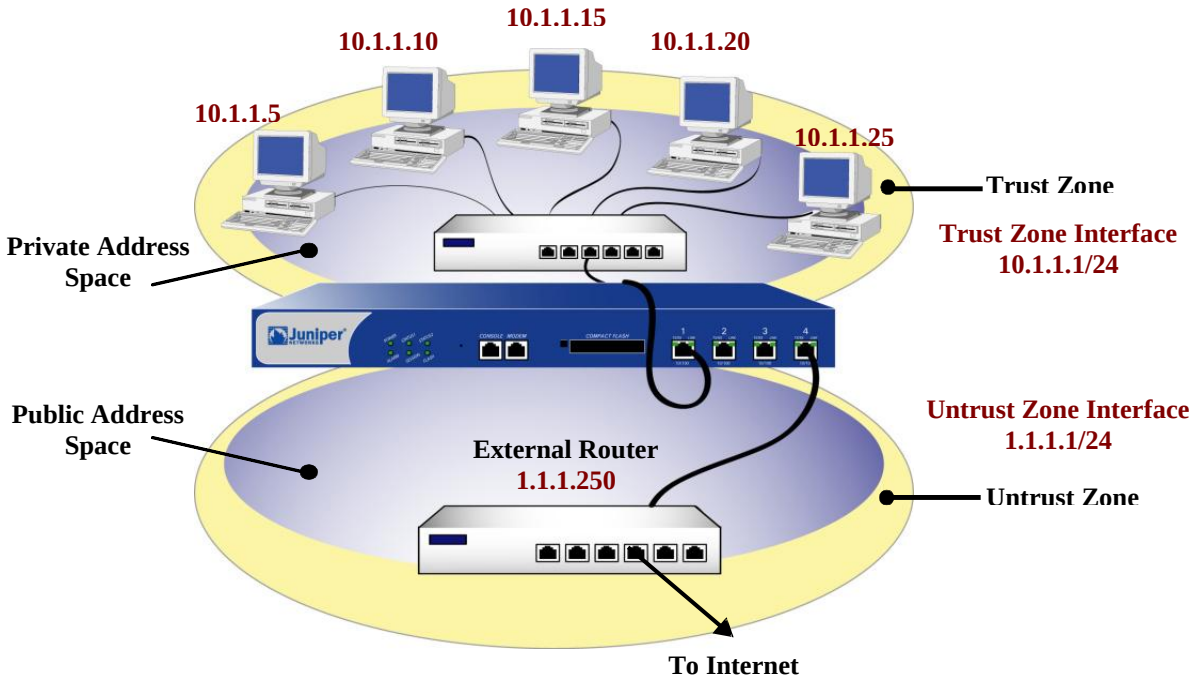

Figure 2.2: NAT Mode

2.3.3 Route Mode

When an interface is in Route mode, the security appliance routes traffic between different zones without performing source NAT (NAT-src); that is, the source address and port number in the IP packet header remain unchanged as it traverses the security appliance. Unlike NAT-src, you do not need to establish mapped IP (MIP) and virtual IP (VIP) addresses to allow inbound traffic to reach hosts when the destination zone interface is in Route mode. Unlike Transparent mode, the interfaces in each zone are on different subnets.

In NAT Mode, Network Address Translation is applied to all IPv4 traffic arriving at the untrust interface. By default, no address translation is provided in Route mode. However, selective network address translation is possible in Route mode using policy definitions. You can determine which traffic to route and on which traffic to perform NAT-src by creating policies that enable NAT-src for specified source addresses on either incoming or outgoing traffic. For network traffic, NAT can use the IPv4 address or addresses of the destination zone interface from a Dynamic IP (DIP) pool, which is in the same subnet as the destination zone interface. For VPN traffic, NAT can use a tunnel interface IPv4 address or an address from its associated DIP pool.

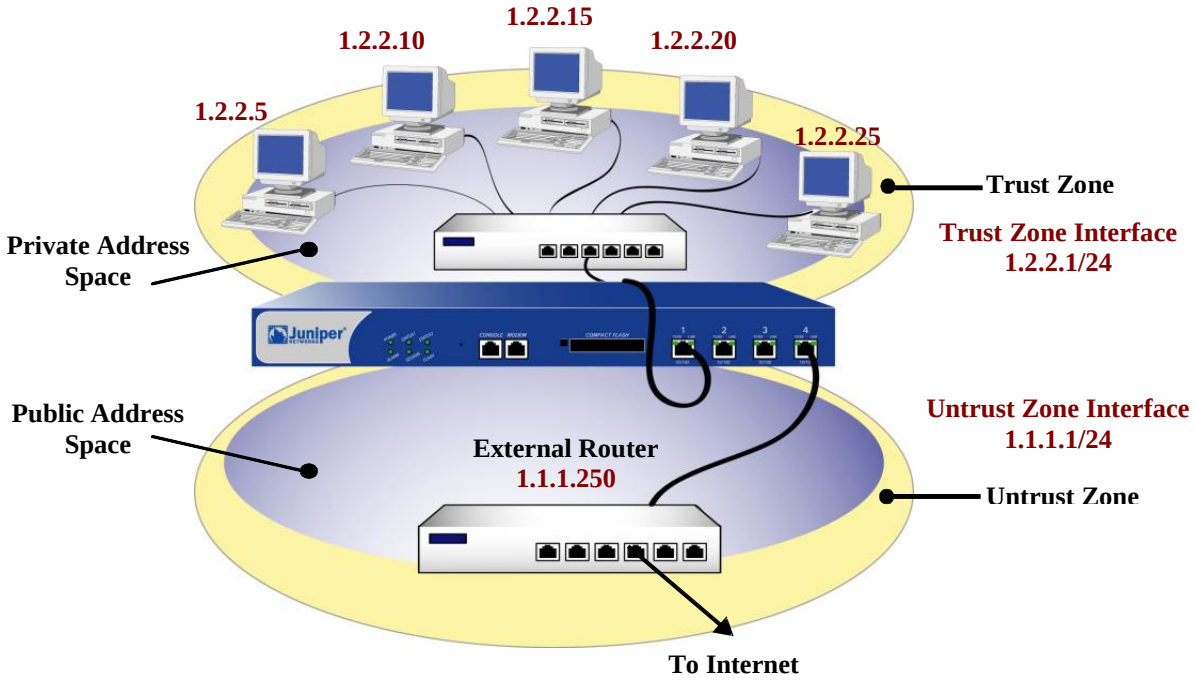

Figure 2.3: Route Mode

2.3.4 Site-to-Site VPN

Site-To-Site VPNs allow an organization to securely connect to a remotely connected network. The TOE supports and defines security claims FDP_IFC.1(1) and FDP_IFF.1(1) for an unauthenticated policy, and FDP_IFC.1(2) and FDP_IFF.1(2) for a VPN policy. The VPN policy utilizes Site-To-Site VPN connections using pre-shared key (PSK) and certificate-based authentication. In order to meet these security functional requirement claims, the TOE must have the appropriate VPN tunnels and permit filters allowing such connectivity and have the appropriate pre-shared key authentication credentials configured. The product supports various methods for VPN connectivity (i.e. Dialup VPN, L2TP VPN, Site-To-Site VPN), authentication (i.e. Manual Key, AutoKey), IPSEC Modes (i.e. Transport, Tunnel), and cryptographic algorithms (i.e. MD5, SHA-1, SHA-256, HMAC, DES, AES). However, the evaluated configuration of the TOE requires that VPN connections are only configured as Site-To-Site VPNs using the IPSEC Tunnel Mode, and any of the following algorithms; SHA-1, SHA-256, HMAC, AES.

While the TOE defines security claims for Site-To-Site VPN connections, an organization is not bound to having VPN configured to meet the evaluated configuration of the TOE. If an organization does not wish to implement the Site-To-Site VPN functionality, then they may exclude it from their configuration of the TOE by ensuring that no VPN tunnels, permit filters, and pre-shared key credentials are established for such connectivity. However in doing so, the organization will not be able to implement the security functionality of the TOE that satisfies the VPN Policy.

The VPN Policy applies to traffic to or from a network interface configured in Transparent Mode that is using a VPN tunnel. The VPN Policy also applies to traffic to or from a network interface configured in Route Mode or NAT Mode that is using a VPN tunnel. The UNAUTHENTICATED Policy applies to traffic to or from a network interface configured in Transparent mode that is not using a VPN tunnel. The UNAUTHENTICATED Policy also applies to traffic to or from a network interface configured in Route Mode or NAT Mode that is not using a VPN tunnel.

2.3.4.1 Policy-Based VPN

Policy-Based VPNs define VPN tunnels through a “tunnel” policy action. A “tunnel” policy action always permits traffic to flow for traffic matching the related routes and services of the VPN tunnel policy.

2.3.4.2 Route-Based VPN

Route-Based VPNs define VPN tunnels using the routing table. For each VPN tunnel, a route is identified to where the VPN tunnel is invoked. Policies can be used in conjunction with the Route-Based VPN to explicitly permit or deny VPN tunnel access based on specified attributes, whereas the Policy-Based VPN only allows the capability to permit specific traffic to a VPN tunnel. Route-Based VPN's are not supported in Transparent mode and only Policy-Based VPN's can be used.

2.4 TOE Architecture

The TOE includes both physical and logical boundaries.

2.4.1 Physical Boundaries

The physical boundary of the security appliances is the physical appliance. The console, which is part of the TOE operational environment, provides the visual I/O for the administrative interface. After the TOE is placed into the evaluated configuration, the administrative interface is provided over an SSH connection using encryption.

The security appliance attaches to physical networks that have been separated into zones through port interfaces.

Security appliances come in several models. Each model differs in the performance capabilities; however all provide the same security functions. Each appliance enforces a security policy for all connection request and traffic flow between any two network zones.

All hardware on which each security appliance operates is part of the TOE. Each security appliance has a custom operating system that is part of the TOE. The operating system, ScreenOS, runs completely in firmware. There is one assumption pertaining to the correct operation of the TOE and that is for the console, which must be a device that can emulate a VT-100 terminal. The console is part of the TOE environment and is expected to correctly display what is sent to it from ScreenOS. Also within the TOE environment are optional servers that can provide time keeping or syslog services. These servers communicate with the TOE over trusted channels using certificate-based authentication and encryption.

The physical boundaries of the security appliance include the interfaces to communicate between an appliance and a network node assigned to a network zone. All network communication flow goes from the sender network node in one zone, through a security appliance, and from a security appliance to the receiving node in another network zone, if the security policy allows the information flow.

2.4.2 Logical Boundaries

This section summarizes the security functions provided by Security Appliances:

- Security audit
- Cryptographic support
- User data protection
- Identification and authentication
- Security management
- Protection of the TSF

2.4.2.1 Security audit

Audit data is stored in memory and is separated into three types of logs; events, traffic logs, and self logs. Events are system-level notifications and alarms which are generated by the system to indicate events such as configuration changes, network attacks detected, or administrators logging in our out of the device. Traffic logs are directly driven by policies that allow traffic to go through the device. Self logs store information on traffic that is dropped and traffic that is sent to the device. Both audit events and traffic messages can be further defined depending on the severity of the message and/or event. Logs are protected and a searching/sorting mechanism of these logs is offered to administrators.

2.4.2.2 Cryptographic support

The Juniper Networks Security Appliances are FIPS 140-2 validated as multi-chip standalone modules. All support the use of AES with SSH using key sizes greater than or equal to 128-bits.

2.4.2.3 User data protection

The user data protection provided by the Security Appliance is provided though the concept of zones. Security policies are applied to the flow of information from network nodes in one zone to network nodes in other zones. These policies control interzone and intrazone information flows.

Traffic from one network node in a zone will only be forwarded to a node in another zone if the connection requests and the traffic satisfy the information flow policies configured in the security appliance. If data is received by an appliance that does not conform to those policies, it will be discarded and an audit record will be sent to the traffic log.

A zone is a logical abstraction on which a security appliance provides services that are typically configurable by the administrator. A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or logical entity that performs a specific function (a function zone).

2.4.2.3.1 Security Zone

A security zone is a segment of network space to which security measures are applied. Multiple security zones can be configured on a single security appliance by sectioning the network into segments to which various security policies may be applied to satisfy the needs of each segment. At a minimum, two security zones must be identified, basically to protect one area of the network from the other. Many security zones can also be established to bring finer granularity to a network security design, without deploying multiple security appliances to do so.

Each security appliance is also configured with a Global Zone. A Global Zone is a security zone without a security zone interface. The Global Zone serves as a storage area for mapped IP (MIP) and virtual IP (VIP) addresses. The predefined Global zone address "Any" applies to all MIPs, VIPs, and other user-defined addresses set in the Global zone. Because traffic going to these addresses is mapped to other addresses, the Global zone does not require an interface for traffic to flow through it.

2.4.2.3.1.1 Security Zone Interface

A security zone interface is an interface in which information can be sent to and from a security zone. Security zones support five types of security zone interfaces, which include physical interfaces, subinterfaces, aggregate interfaces, redundant interfaces, and virtual security interfaces. However, the evaluated configuration of the TOE may only utilize the physical interfaces, aggregate interfaces, and redundant interfaces.

2.4.2.3.1.1.1 Physical Interface

Each physical network port on the security appliance represents a physical interface, and the name of the interface is predefined. The name of a physical interface is composed of the media type, slot number (for some security appliances), and port number, for example, ethernet3/2 or ethernet2. A physical interface can bind to any security zone where it acts as a doorway through which traffic enters and exits the zone. Without a physical interface, no traffic can access the zone or leave it.

2.4.2.3.1.1.2 Aggregate Interface

The Juniper Networks NetScreen-5000 series supports aggregate interfaces. An aggregate interface is the accumulation of two or more physical interfaces, each of which shares the traffic load directed to the IP address of the aggregate interface equally among them. By using an aggregate interface, the amount of bandwidth available to a single IP address can be increased. Also, if one member of an aggregate interface fails, the other member or members can continue processing traffic, although with less bandwidth than previously available.

2.4.2.3.1.1.3 Redundant Interface

A redundant interface consists of binding two physical interfaces together to create one redundant interface, which you can then bind to a security zone. One of the two physical interfaces acts as the primary interface and handles all the traffic directed to the redundant interface. The other physical interface acts as the secondary interface and stands by in case the active interface experiences a failure. If that occurs, traffic to the redundant interface fails over to the secondary interface, which becomes the new primary interface. The use of redundant interfaces provides a first line of redundancy before escalating a failover to the device level.

2.4.2.3.2 Tunnel Zone

The concept of a tunnel zone is described only as background to for policy enforcement based upon zones.

A tunnel zone is a logical segment that hosts one or more tunnel interfaces. A tunnel zone is conceptually affiliated with a security zone in a "child-parent" relationship. The security zone acting as the "parent" provides the firewall protection to the encapsulated traffic. The tunnel zone provides packet encapsulation/decapsulation, and by supporting tunnel interfaces with IP addresses and net masks that can host mapped IP (MIP) addresses and dynamic IP (DIP) pools, can also provide policy-based NAT services. The security appliance uses the routing information for the carrier zone to direct traffic to the tunnel endpoint. The default tunnel zone is Untrust-Tun, and it is associated with the Untrust zone. Other tunnel zones can be created and bound to other security zones, with a maximum of one tunnel zone per carrier zone per virtual system. Virtual systems, however, are outside the scope of the evaluated configuration.

2.4.2.3.3 Function Zone

The function zone is a zone that performs a specific function. Functional zones support five types of zones, which include null zones, MGT zones, HA zones, self zones, and VLAN zones. However, the evaluated configuration of the TOE may only utilize the null zones and self zones. Each zone exists for a single purpose, as explained below.

2.4.2.3.3.1 Null Zone

This zone serves as temporary storage for any interfaces that are not bound to any other zone.

2.4.2.3.3.2 Self Zone

This zone hosts the interface for remote management connections. When connecting to the security appliance via HTTP, SSH, or Telnet, the self zone is used. Remote management is supported in the evaluated configuration of the TOE only via SSH (the other potential management interfaces must be disabled).

2.4.2.4 Identification and authentication

The security appliances provide an authentication mechanism for administrative users through an internal authentication database. Administrative login is supported through the locally connected console for initial configuration, or remotely via an SSH protected communication channel. The TOE operates in a mode that has been certified to FIPS 140-2 level 2 overall, and supports AES encryption for the SSH protected communication channel.

A known administrator user id and its corresponding authentication data must be entered correctly in order for the administrator to successfully logon and thereafter gain access to administrative functions. For local authentication, all administrator user name and password pairs are managed in a database internal to the security appliance. Excessive failed login attempts while initiating a remote administration session can cause the session being created to be closed.

2.4.2.5 Security management

Every security appliance provides a command line administrative interface and supports remote administration through an SSH command line interface. SSH provides for the protection of remote administration activity from

both disclosure and modification. Neither the web interface nor the Network and Security Manager[3] are part of the evaluated configuration.

To execute the CLI, the administrator can establish a trusted SSH connection to the security appliance. The authorized administrator must be successfully identified and authenticated before they are permitted to perform any security management functions on the TOE.

The Security Appliances also support distinct administrative roles: Root Administrator, Audit Administrator, Cryptographic Administrator and Security Administrator. In addition to these administrative roles, an administrator may be given a read-write or read-only attribute that affects that administrator's ability to change the device's configuration data. All of these roles are considered to be authorized administrators.

More details about these management operations available to administrators can be found in Section 6.1.5, 'Security management'.

2.4.2.6 Protection of the TSF

Each security appliance is a hardware and firmware device that protects itself largely by offering only a minimal logical interface to the network and attached nodes. ScreenOS is a special purpose OS that provides no general purpose programming capability. All network traffic from one network zone to another or between two networks within the same network zone passes through the TOE; however, no protocol services are provided for non-administrative user communication with the security appliance itself. The TOE also utilizes a hardware clock to maintain and provide reliable time stamps.

The TOE can be configured by an administrator to use the Network Time Protocol (NTP), to synchronize the hardware clock with an external time server.

2.5 TOE Documentation

Juniper Networks offers a series of documents that describe the installation of Security Appliances as well as guidance for subsequent use and administration of the applicable security features.

- ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide, Volume 1: Overview
- ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide, Volume 2: Fundamentals
- ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide, Volume 3: Administration
- ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide, Volume 4: Attack Detection
- ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide, Volume 5: VPNs
- ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide, Volume 8: Address Translation
- ScreenOS CLI Reference Guide: IPv4 Command Descriptions
- ScreenOS 6.3.0 Message Log Reference Guide
- Juniper Networks ScreenOS 6.3 Evaluated Configuration for Common Criteria, EAL4
- SSG 5 Hardware Installation and Configuration Guide
- SSG 20 Hardware Installation and Configuration Guide
- SSG 140 Hardware Installation and Configuration Guide
- SSG 300M-series Hardware Installation and Configuration Guide
- SSG 500M-series Hardware Installation and Configuration Guide
- ISG 1000 Hardware Installation and Configuration Guide
- ISG 2000 Hardware Installation and Configuration Guide
- NetScreen-5000 Series Hardware Installation and Configuration Guide

Note: Several sections of the ScreenOS 6.3.0 Concepts and Example, ScreenOS Reference Guide are NOT included as part of the TOE documentation. These sections were excluded because this ST makes no claims regarding the functionality within these sections. Operation of the TOE with these features is not part of this evaluation.

[3] Network and Security Manager is Juniper Network's enterprise-level management software application.

3 Security Problem Definition

The security problem definition has been drawn from a validated PP (TFFW BR). Please consult that PP for the description of the security environment. The policies, threats and assumptions from those PPs have been copied here for convenience. However, the TFFW BR PP contains the definitive statement of the security problem definition.

3.1 Threats

The following threats are addressed either by the TOE or the environment. The threats discussed below are addressed by Protection Profile-compliant TOEs. The threat agents are either unauthorized persons or external IT entities not authorized to use the TOE itself.

T.NOAUTH	An unauthorized person may attempt to bypass the security of the TOE so as to access and use security functions and/or non-security functions provided by the TOE.
T.REPEAT	An unauthorized person may repeatedly try to guess authentication data in order to use this information to launch attacks on the TOE.
T.REPLAY	An unauthorized person may use valid identification and authentication data obtained to access functions provided by the TOE.
T.ASPOOF	An unauthorized person may carry out spoofing in which information flow through the TOE into a connected network by using a spoofed source address.
T.MEDIAT	An unauthorized person may send impermissible information through the TOE which results in the exploitation of resources on the internal network.
T.OLDINF	Because of a flaw in the TOE functioning, an unauthorized person may gather residual information from a previous information flow or internal TOE data by monitoring the padding of the information flows from the TOE.
T.PROCOM	An unauthorized person or unauthorized external IT entity may be able to view, modify, and/or delete security related information that is sent between a remotely located authorized administrator and the TOE
T.AUDACC	Persons may not be accountable for the actions that they conduct because the audit records are not reviewed, thus allowing an attacker to escape detection.
T.SELPRO	An unauthorized person may read, modify, or destroy security critical TOE configuration data.
T.AUDFUL	An unauthorized person may cause audit records to be lost or prevent future records from being recorded by taking actions to exhaust audit storage capacity, thus masking an attackers actions.

3.2 Assumptions

The following conditions are assumed to exist in the operational environment.

A.PHYSEC	The TOE is physically secure.
A.LOWEXP	The threat of malicious attacks aimed at discovering exploitable vulnerabilities is considered low.

A.GENPUR	There are no general-purpose computing capabilities (e.g., the ability to execute arbitrary code or applications) and storage repository capabilities on the TOE.
A.PUBLIC	The TOE does not host public data.
A.NOEVIL	Authorized administrators are non-hostile and follow all administrator guidance; however, they are capable of error.
A.SINGEN	Information can not flow among the internal and external networks unless it passes through the TOE.
A.DIRECT	Human users within the physically secure boundary protecting the TOE may attempt to access the TOE from some direct connection (e.g., a console port) if the connection is part of the TOE.
A.NOREMO	Human users who are not authorized administrators can not access the TOE remotely from the internal or external networks.
A.REMACC	Authorized administrators may access the TOE remotely from the internal and external networks.

3.3 Policies

P.INTEGRITY The TOE shall support the IETF Internet Protocol Security Encapsulating Security Payload (IPSEC ESP) as specified in RFC 2406. Sensitive information transmitted to a peer TOE shall apply integrity mechanisms as specified in Use of HMAC-SHA-1-96 within ESP and AH (RFC 2404)

4 Security Objectives

All of the Security Objectives have been drawn from a validated PP (TFFW BR). Please consult that PP for a description of the security objectives. The security objectives from that PP have been copied here for convenience. However, the TFFW BR PP contains the definitive statement of security objectives.

4.1 Security Objectives for the TOE

The following are the IT security objectives for the TOE:

O.INTEGRITY The TOE must be able to protect the integritty of data transmitted to a peer via encryption and provide IPSec authentication for such data. Upon receipt of data from a peer TOE, the TOE must be able to decrypt the data and verify that the received data accurately represents the data that was originally transmitted.

O.IDAUTH The TOE must uniquely identify and authenticate the claimed identity of all users, before granting a user access to TOE functions.

O.SINUSE The TOE must prevent the reuse of authentication data for users attempting to authenticate at the TOE from a connected network.

O.MEDIAT The TOE must mediate the flow of all information from users on a connected network to users on another connected network, and must ensure that residual information from a previous information flow is not transmitted in any way.

O.SECSTA Upon initial start-up of the TOE or recovery from an interruption in TOE service, the TOE must not compromise its resources or those of any connected network.

O.ENCRYP The TOE must protect the confidentiality of its dialogue with an authorized administrator through encryption, if the TOE allows administration to occur remotely from a connected network.

O.SELPRO The TOE must protect itself against attempts by unauthorized users to bypass, deactivate, or tamper with TOE security functions.

O.AUDREC The TOE must provide a means to record a readable audit trail of security-related events, with accurate dates and times, and a means to search and sort the audit trail based on relevant attributes.

O.ACCOUN The TOE must provide user accountability for information flows through the TOE and for authorized administrator use of security functions related to audit.

O.SECFUN The TOE must provide functionality that enables an authorized administrator to use the TOE security functions, and must ensure that only authorized administrators are able to access such functionality.

O.LIMEXT The TOE must provide the means for an authorized administrator to control and limit access to TOE security functions by an authorized external IT entity.

4.2 Security Objectives for the Environment

All of the assumptions stated in section 3.2 are considered to be security objectives for the operational environment. The following are the PP non-IT security objectives, which, in addition to those assumptions, are to be satisfied without imposing technical requirements on the TOE. That is, they will not require the implementation of functions in the TOE hardware and/or software. Thus, they will be satisfied largely through application of procedural or administrative measures.

A.PHYSEC	The TOE is physically secure.
A.LOWEXP	The threat of malicious attacks aimed at discovering exploitable vulnerabilities is considered low.
A.GENPUR	There are no general-purpose computing capabilities (e.g., the ability to execute arbitrary code or applications) and storage repository capabilities on the TOE.
A.PUBLIC	The TOE does not host public data.
A.NOEVIL	Authorized administrators are non-hostile and follow all administrator guidance; however, they are capable of error.
A.SINGEN	Information can not flow among the internal and external networks unless it passes through the TOE.
A.DIRECT	Human users within the physically secure boundary protecting the TOE may attempt to access the TOE from some direct connection (e.g., a console port) if the connection is part of the TOE.
A.NOREMO	Human users who are not authorized administrators can not access the TOE remotely from the internal or external networks.
A.REMACC	Authorized administrators may access the TOE remotely from the internal and external networks.
O.GUIDAN	The TOE must be delivered, installed, administered, and operated in a manner that maintains security.
O.ADMTRA	Authorized administrators are trained as to establishment and maintenance of security policies and practices.

5 IT Security Requirements

The security requirements for the TOE have all been drawn from Parts 2 and 3 of the Common Criteria as well as from the TFFW BR PP. The security functional requirements have been selected to correspond to the actual security functions implemented by the TOE while the assurance requirements have been selected to offer assurance that those security functions are properly realized.

5.1 TOE Security Functional Requirements

The following table describes the SFRs that are satisfied by Security Appliances.

Requirement Class	Requirement Component
FAU: Security Audit	FAU_GEN.1: Audit data generation
	FAU_SAR.1: Audit review
	FAU_SAR.3: Selectable audit review
	FAU_STG.1: Protected audit trail storage
	FAU_STG.4: Prevention of audit data loss
FCS: Cryptographic Support	FCS_COP.1(1): Cryptographic Operation (remote sessions)
	FCS_COP.1(2): Cryptographic Operation (VPN operations)
FDP: User Data Protection	FDP_IFC.1(1): Subset information flow control (Unauthenticated Policy)
	FDP_IFC.1(2): Subset information flow control (VPN Policy)
	FDP_IFF.1(1): Simple security attributes (Unauthenticated Policy)
	FDP_IFF.1(2): Simple security attributes (VPN Policy)
FIA: Identification and Authentication	FIA_AFL.1: Authentication failure handling
	FIA_ATD.1: User attribute definition
	FIA_UAU.1: Timing of authentication
	FIA_UID.2: User identification before any action
	FIA_UAU.4: Single-use authentication mechanisms
FMT: Security Management	FMT_MOF.1: Management of security functions behavior
	FMT_MSA.3: Static attribute initialization
	FMT_SMR.1: Security Roles
	FMT_SMF.1: Specification of Management Functions
FPT: TSF Protection	FPT_STM.1: Reliable time stamps

Table 5-1 TOE Security Functional Requirements

5.1.1 Security audit (FAU)

5.1.1.1 FAU_GEN.1: Audit data generation

FAU_GEN.1.1 The TSF shall be able to generate an audit record of the following auditable events:

a) Start-up and shutdown of the audit functions;

b) All relevant auditable events for the minimal or basic level of audit specified in Table 5-2; and

c) [the event in Table 5-2 listed at the "extended" level].

FAU_GEN.1.2 The TSF shall record within each audit record at least the following information:

a) Date and time of the event, type of event, subjects identities, outcome (success or failure) of the event; and

b) For each audit event type, based on the auditable event definitions of the functional components included in the PP/ST, [information specified in column four of Table 5-2].

c)

Functional Component	Level	Auditable Event	Additional Audit Record Contents
FMT_SMR.1	Minimal	Modifications to the group of users that are part of the authorized administrator role	The identity of the authorized administrator performing the modification and the user identity being associated with the authorized administrator role
FIA_UID.2	Basic	All use of the user identification mechanism	The user identities provided to the TOE
FIA_UAU.1	Basic	Any use of the authentication mechanism	The user identities provided to the TOE
FIA_AFL.1	Minimal	The reaching of the threshold for unsuccessful authentication attempts and the subsequent restoration by the authorized administrator of the users capability to authenticate	The identity of the offending user and the authorized administrator
FDP__IFF.1(1)	Basic	All decisions on requests for information flow	The presumed addresses of the source and destination subject
FDP_IFF.1(2)	Basic	All decisions to permit and deny flows	The presumed addresses of the source and destination subject
FCS_COP.1	Minimal	Success and failure, and the type of cryptographic operation	The identity of the external IT entity attempting to perform the cryptographic operation
FPT_STM.1	Minimal	Changes to the time	The identity of the authorized administrator performing the operation
FMT_MOF.1	Extended	Use of the functions listed in this requirement pertaining to audit	The identity of the authorized administrator performing the operation

Table 5-2 Auditable Events

Note: FAU_GEN.1.1b is copied directly from the PP. It appears that the PP was attempting to indicate that the required events were those listed in the table in the PP. This ST is simply repeating the PP and has reproduced the table from the PP. Audit events for SFRs included in the PP but not listed in the table were omitted by the PP author and thus are omitted from this ST.

5.1.1.2 FAU_SAR.1: Audit review

FAU_SAR.1.1 The TSF shall provide [an authorized administrator] with the capability to read [all audit trail data] from the audit records.

FAU_SAR.1.2 The TSF shall provide the audit records in a manner suitable for the user to interpret the information.

5.1.1.3 FAU_SAR.3: Selectable audit review

FAU_SAR.3.1 The TSF shall provide the ability to perform searches and sorting of audit data based on:

a) [presumed subject address;
b) ranges of dates;
c) ranges of times;
d) ranges of addresses]

5.1.1.4 FAU_STG.1: Protected audit trail storage

FAU_STG.1.1 **The TSF shall protect the stored audit records in the audit trail from unauthorised deletion.**.

FAU_STG.1.2 The TSF shall be able to prevent unauthorized modifications to the audit records in the audit trail.

5.1.1.5 FAU_STG.4: Prevention of audit data loss

FAU_STG.4.1 The TSF shall *prevent auditable events, except those taken by the authorized administrator* and [shall limit the number of audit records lost] if the audit trail is full.

5.1.2 Cryptographic support (FCS)

5.1.2.1 FCS_COP.1(1): Cryptographic Operation (remote sessions)

FCS_COP.1.1(1) The TSF shall perform [encryption of remote authorized administrator sessions] in accordance with a specified cryptographic algorithm:

- [AES (Advanced Encryption Standard as specified in FIPS 197) encryption (as specified in SP 800-67)] and cryptographic key sizes [that are at least 128 binary digits in length] that meet the following: [FIPS PUB 140-2 (Level 1)].

Application Note from BRPP: This requirement is applicable only if the TOE includes the capability for the authorized administrator to perform security functions remotely from a connected network. In this case, AES encryption must protect the communications between the authorized administrator and the TOE, and the associated cryptographic module(s) must comply at a minimum with FIPS PUB 140-2 Level 1.

5.1.2.2 FCS_COP.1(2): Cryptographic Operation (VPN operations)

FCS_COP.1.1(2) The TSF shall perform [**encryption/decryption on VPN traffic per the VPN SFP**] in accordance with a specified cryptographic algorithm: [**see table Table 5-3**] and cryptographic key sizes [**see table Table 5-3**] that meet the following: [FIPS PUB 140-2 (Level 1)].

Table 5-3 FCS_COP.1(2) Algorithms and Key Sizes

Algorithm	Key Size
AES (Advanced Encryption Standard as specified in FIPS 197) encryption (as specified in SP 800-67)	**at least 128 binary digits in length**
ECDSA (Elliptic Curve Digital Signature Algorithm per FIPS 186-3) using only the NIST curve(s) P-256	**at least 256 binary digits in length**
SHA-256 (Secure Hash Algorithm per FIPS 180-2)	**256 binary digits in length**

5.1.3 User data protection (FDP)

5.1.3.1 FDP_IFC.1(1): Subset information flow control (Unauthenticated Policy)

FDP_IFC.1.1(1) The TSF shall enforce the [UNAUTHENTICATED SFP] on:

a) [subjects: unauthenticated external IT entities that send and receive information through the TOE to one another;

b) information: traffic sent through the TOE from one subject to another;

c) operation: pass information].

5.1.3.2 FDP_IFC.1(2): Subset information flow control (VPN Policy)

FDP_IFC.1.1(2) The TSF shall enforce the **[VPN SFP]** on **[**

a) Subjects:

a. source subject: TOE interface on which information is received;

b. destination subject: TOE interface to which information is destined.

b) information: network packets; and

a) operations:

i. pass packets without modifying;

ii. send IPSEC encrypted and authenticated packets to a peer TOE using ESP in tunnel mode as defined in RFC 2406;

iii. decrypt, verify authentication and pass received packets from a peer TOE in tunnel mode using ESP;

iv. drop packets from passing].

5.1.3.3 FDP_IFF.1(1): Simple security attributes (Unauthenticated Policy)

FDP_IFF.1.1(1) The TSF shall enforce the [UNAUTHENTICATED SFP] based on at least the following types of subject and information security attributes:

a) [subject security attributes:

- presumed address;
- **[none]**;

b) information security attributes:

- presumed address of source subject;
- presumed address of destination subject;
- transport layer protocol;
- TOE interface on which traffic arrives and departs;
- service;
- **[none]**].

FDP_IFF.1.2(1) The TSF shall permit an information flow between a controlled subject and another controlled subject via a controlled operation if the following rules hold:

a) [Subjects on an internal network can cause information to flow through the TOE to another connected network if:

- all the information security attribute values are unambiguously permitted by the information flow security policy rules, where such rules may be composed from all possible combinations of the values of the information flow security attributes, created by the authorized administrator;
- the presumed address of the source subject, in the information, translates to an internal network address;
- and the presumed address of the destination subject, in the information, translates to an address on the other connected network.

b) Subjects on the external network can cause information to flow through the TOE to another connected network if:

- all the information security attribute values are unambiguously permitted by the information flow security policy rules, where such rules may be composed from all possible combinations of the values of the information flow security attributes, created by the authorized administrator;

- the presumed address of the source subject, in the information, translates to an external network address;
- and the presumed address of the destination subject, in the information, translates to an address on the other connected network.]

FDP_IFF.1.3(1) The TSF shall enforce the [none].

FDP_IFF.1.4(1) The TSF shall provide the following [none].

FDP_IFF.1.5(1) The TSF shall explicitly authorize an information flow based on the following rules: [none].

FDP_IFF.1.6(1) The TSF shall explicitly deny an information flow based on the following rules:

a) [The TOE shall reject requests for access or services where the information arrives on an external TOE interface, and the presumed address of the source subject is an external IT entity on an internal network;

b) The TOE shall reject requests for access or services where the information arrives on an internal TOE interface, and the presumed address of the source subject is an external IT entity on the external network;

c) The TOE shall reject requests for access or services where the information arrives on either an internal or external TOE interface, and the presumed address of the source subject is an external IT entity on a broadcast network;

d) The TOE shall reject requests for access or services where the information arrives on either an internal or external TOE interface, and the presumed address of the source subject is an external IT entity on the loopback network.]

5.1.3.4 FDP_IFF.1(2): Simple security attributes (VPN Policy)

FDP_IFF.1.1(2) The TSF shall enforce the [**VPN SFP**] based on the following types of subject and information security attributes:

a) [**Source subject security attributes:**
- **set of source subject identifiers**

b) **Destination subject security attributes:**
- **Set of destination subject identifiers;**

c) **Information security attributes:**
- **presumed identity of source subject;**
- **identity of destination subject**;]

FDP_IFF.1.2(2) The TSF shall permit an information flow between a source subject and a destination subject via a controlled operation if the following rules hold:

- [**the presumed identity of the source subject is in the set of source subject identifiers;**
- **the identity of the destination subject is in the set of source destination identifiers;**
- **the information security attributes match the attributes in an information flow policy rule (contained in the information flow policy ruleset defined by the Security Administrator) according to the following algorithm**

For ROUTE Mode operation:
- **Packets with source and destination addresses in different zones the Interzone set of rules is applied, the Global set of rules is applied and if a match was not found the default action is performed.**
- **Packets with source and destination addresses in the same zone the Intrazone set of rules is applied, the Global set of rules is applied and if a match was not found the default action is performed.**

For Transparentt Mode operation:
- **Packets with source and destination addresses in different zones the Interzone set of rules is applied and if a match was not found the default action is performed.**

For all modes:

- **Evaluation of a set of rules is done by evaluating rules sequentially within the set of rules searching for the first matching rule and performing the action specified by that rule to the packet**
- **A rule matches if all of the information security attributes are unambiguously permitted by the rule; and**

▪ **the selected information flow policy rule specifies that the information flow is to be permitted, and what specific operation from FDP_IFC.1(2) is to be applied to that information flow**].

FDP_IFF.1.3(1) The TSF shall enforce the **[none]**.

FDP_IFF.1.4(1) The TSF shall provide the following **[none]**.

FDP_IFF.1.5(1) The TSF shall explicitly authorize an information flow based on the following rules: **[none]**.

FDP_IFF.1.6(2) The TSF shall explicitly deny an information flow based on the following rules:

a) [**The TOE shall reject requests for access or services where the presumed source identity of the information received by the TOE is not included in the set of source identifiers for the source subject;**

b) **The TOE shall reject requests for access or services where the presumed source identity of the information received by the TOE specifies a broadcast identity;**

c) **The TOE shall reject requests for access or services where the presumed source identity of the information received by the TOE specifies a loopback identifier;**

d) **The TOE shall reject requests in which the information received by the TOE contains the route (set of host network identifiers) by which information shall flow from the source subject to the destination subject.)**].

5.1.3.5 FDP_RIP.2: Full residual information protection

FDP_RIP.2.1 The TSF shall ensure that any previous information content of a resource is made unavailable upon the [***allocation of the resource from***] all objects.

5.1.4 Identification and authentication (FIA)

5.1.4.1 FIA_AFL.1: Authentication failure handling

FIA_AFL.1.1 The TSF shall detect when [**a Security Administrator-configurable positive integer**] of unsuccessful authentication attempts occur related to [external IT entities attempting to authenticate from an internal or external network.]

FIA_AFL.1.2 When the defined number of unsuccessful authentication attempts has been met or surpassed, the TSF shall [prevent the offending external IT entity from successfully authenticating until an authorized administrator takes some action to make authentication possible for the external IT entity in question.]

5.1.4.2 FIA_ATD.1: User attribute definition

FIA_ATD.1.1 The TSF shall maintain the following list of security attributes belonging to individual users:

a) [identity;

b) association of a human user with the authorized administrator role;

c) **[password]**] .

5.1.4.3 FIA_UAU.1: Timing of authentication

FIA_UAU.1.1 The TSF shall allow [identification as stated in FIA_UID.2] on behalf of the authorized administrator or authorized external IT entity accessing the TOE to be performed before the authorized administrator or authorized external IT entity is authenticated.

FIA_UAU.1.2 The TSF shall require each authorized administrator or authorized external IT entity to be successfully authenticated before allowing any other TSF-mediated actions on behalf of that authorized administrator or authorized IT entity.

Note: The BRPP indicates 'The TOE is permitted to pass information before users are authenticated.' As such, it should be understood that this SFR pertains only to administrators or other network entities obtaining authorized access to TOE services and not to network traffic flowing in accordance with the UNAUTHENTICATED SFP.

5.1.4.4 FIA_UID.2: User identification before any action

FIA_UID.2.1 The TSF shall require each user to identify itself before allowing any other TSF-mediated actions on behalf of that user.

5.1.4.5 FIA_UAU.4: Single-use authentication mechanisms

FIA_UAU.4.1 The TSF shall prevent reuse of authentication data related to [authentication attempts from either an internal or external network by:

a) authorized administrators;

b) authorized external IT entities].

5.1.5 Security management (FMT)

5.1.5.1 FMT_MOF.1: Management of security functions behavior

FMT_MOF.1.1 The TSF shall restrict the ability to determine and modify the behavior of the functions

a) [start-up and shutdown;

b) create, delete, modify, and view information flow security policy rules that permit or deny information flows;

c) create, delete, modify, and view user attribute values defined in FIA_ATD.1;

d) enable and disable single-use authentication mechanisms in FIA_UAU.4 (if the TOE supports authorized IT entities and/or remote administration from either an internal or external network);

e) modify and set the threshold for the number of permitted authentication attempt failures (if the TOE supports authorized IT entities and/or remote administration from either an internal or external network);

f) restore authentication capabilities for users that have met or exceeded the threshold for permitted authentication attempt failures (if the TOE supports authorized IT entities and/or remote administration from either an internal or external network);

g) enable and disable external IT entities from communicating to the TOE (if the TOE supports authorized external IT entities);

h) modify and set the time and date;

i) archive, create, delete, empty, and review the audit trail;

j) backup of user attribute values, information flow security policy rules, and audit trail data,;

k) recover to the state following the last backup;

l) additionally, if the TSF supports remote administration from either an internal or external network:

- enable and disable remote administration from internal and external networks;
- restrict addresses from which remote administration can be performed;

m) **[execute TSF self tests]**]

to [an authorized administrator].

5.1.5.2 FMT_MSA.3: Static attribute initialization

FMT_MSA.3.1 The TSF shall enforce the [UNAUTHENTICATED SFP and VPN SFP] to provide restrictive default values for the information flow security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The TSF shall allow the [authorized administrator] to specify alternative initial values to override the default values when an object or information is created.

5.1.5.3 FMT_SMF.1: Specification of Management Functions

FMT_SMF.1.1 The TSF shall be capable of performing the following security management functions: [

1. **start-up and shutdown;**
2. **create, delete, modify, and view information flow security policy rules that permit or deny information flows;**
3. **create, delete, modify, and view user attribute values defined in FIA_ATD.1;**
4. **enable and disable single-use authentication mechanisms in FIA_UAU.4 (if the TOE supports authorized IT entities and/or remote administration from either an internal or external network);**
5. **modify and set the threshold for the number of permitted authentication attempt failures (if the TOE supports authorized IT entities and/or remote administration from either an internal or external network);**
6. **restore authentication capabilities for users that have met or exceeded the threshold for permitted authentication attempt failures (if the TOE supports authorized IT entities and/or remote administration from either an internal or external network);**
7. **enable and disable external IT entities from communicating to the TOE (if the TOE supports authorized external IT entities);**
8. **modify and set the time and date;**
9. **archive, create, delete, empty, and review the audit trail;**
10. **backup of user attribute values, information flow security policy rules, and audit trail data;**
11. **recover to the state following the last backup;**
12. **additionally, if the TSF supports remote administration from either an internal or external network:**
13. **enable and disable remote administration from internal and external networks;**
14. **restrict addresses from which remote administration can be performed**].

5.1.5.4 FMT_SMR.1: Security Roles

FMT_SMR.1.1 The TSF shall maintain the roles: [authorized administrator].
FMT_SMR.1.2 The TSF shall be able to associate human users with the authorized administrator role.

5.1.6 Protection of the TSF (FPT)

5.1.6.1 FPT_STM.1: Reliable time stamps

FPT_STM.1.1 The TSF shall be able to provide reliable time stamps for its own use.

5.2 TOE Security Assurance Requirements

The security assurance requirements for the TOE are the EAL 2 augmented with ALC_FLR.2 components as specified in Part 3 of the Common Criteria. No operations are applied to the assurance components.

Requirement Class	Requirement Component
ADV: Development	ADV_ARC.1: Security architecture description
	ADV_FSP.2: Security-enforcing functional specification
	ADV_TDS.1: Basic design
AGD: Guidance documents	AGD_OPE.1: Operational user guidance
	AGD_PRE.1: Preparative procedures
ALC: Life-cycle support	ALC_CMC.2: Use of a CM system
	ALC_CMS.2: Parts of the TOE CM coverage
	ALC_DEL.1: Delivery procedures
	ALC_FLR.2: Flaw reporting procedures
ATE: Tests	ATE_COV.1: Evidence of coverage
	ATE_FUN.1: Functional testing
	ATE_IND.2: Independent testing - sample
AVA: Vulnerability assessment	AVA_VAN.2: Vulnerability analysis

Table 5-4 Basic Robustness Assurance Requirements

5.2.1 Development (ADV)

5.2.1.1 Security architecture description (ADV_ARC.1)

ADV_ARC.1.1d The developer shall design and implement the TOE so that the security features of the TSF cannot be bypassed.

ADV_ARC.1.2d The developer shall design and implement the TSF so that it is able to protect itself from tampering by untrusted active entities.

ADV_ARC.1.3d The developer shall provide a security architecture description of the TSF.

ADV_ARC.1.1c The security architecture description shall be at a level of detail commensurate with the description of the SFR-enforcing abstractions described in the TOE design document.

ADV_ARC.1.2c The security architecture description shall describe the security domains maintained by the TSF consistently with the SFRs.

ADV_ARC.1.3c The security architecture description shall describe how the TSF initialisation process is secure.

ADV_ARC.1.4c The security architecture description shall demonstrate that the TSF protects itself from tampering.

ADV_ARC.1.5c The security architecture description shall demonstrate that the TSF prevents bypass of the SFR-enforcing functionality.

ADV_ARC.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.1.2 Security-enforcing functional specification (ADV_FSP.2)

ADV_FSP.2.1d The developer shall provide a functional specification.

ADV_FSP.2.2d The developer shall provide a tracing from the functional specification to the SFRs.

ADV_FSP.2.1c The functional specification shall completely represent the TSF.

ADV_FSP.2.2c The functional specification shall describe the purpose and method of use for all TSFI.

ADV_FSP.2.3c The functional specification shall identify and describe all parameters associated with each TSFI.

ADV_FSP.2.4c For SFR-enforcing TSFIs, the functional specification shall describe the SFR-enforcing actions associated with the TSFI.

ADV_FSP.2.5c For SFR-enforcing TSFIs, the functional specification shall describe direct error messages resulting from processing associated with the SFR-enforcing actions.

ADV_FSP.2.6c The tracing shall demonstrate that the SFRs trace to TSFIs in the functional specification.

ADV_FSP.2.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

ADV_FSP.2.2e The evaluator shall determine that the functional specification is an accurate and complete instantiation of the SFRs.

5.2.1.3 Basic design (ADV_TDS.1)

ADV_TDS.1.1d The developer shall provide the design of the TOE.

ADV_TDS.1.2d The developer shall provide a mapping from the TSFI of the functional specification to the lowest level of decomposition available in the TOE design.

ADV_TDS.1.1c The design shall describe the structure of the TOE in terms of subsystems.

ADV_TDS.1.2c The design shall identify all subsystems of the TSF.

ADV_TDS.1.3c The design shall describe the behaviour of each SFR-supporting or SFR-non-interfering TSF subsystem in sufficient detail to determine that it is not SFR-enforcing.

ADV_TDS.1.4c The design shall summarise the SFR-enforcing behaviour of the SFR-enforcing subsystems.

ADV_TDS.1.5c The design shall provide a description of the interactions among SFR-enforcing subsystems of the TSF, and between the SFR-enforcing subsystems of the TSF and other subsystems of the TSF.

ADV_TDS.1.6c The mapping shall demonstrate that all behaviour described in the TOE design is mapped to the TSFIs that invoke it.

ADV_TDS.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

ADV_TDS.1.2e The evaluator shall determine that the design is an accurate and complete instantiation of all security functional requirements.

5.2.2 Guidance documents (AGD)

5.2.2.1 Operational user guidance (AGD_OPE.1)

AGD_OPE.1.1d The developer shall provide operational user guidance.

AGD_OPE.1.1c The operational user guidance shall describe, for each user role, the user-accessible functions and privileges that should be controlled in a secure processing environment, including appropriate warnings.

AGD_OPE.1.2c The operational user guidance shall describe, for each user role, how to use the available interfaces provided by the TOE in a secure manner.

AGD_OPE.1.3c The operational user guidance shall describe, for each user role, the available functions and interfaces, in particular all security parameters under the control of the user, indicating secure values as appropriate.

AGD_OPE.1.4c The operational user guidance shall, for each user role, clearly present each type of security-relevant event relative to the user-accessible functions that need to be performed, including changing the security characteristics of entities under the control of the TSF.

AGD_OPE.1.5c The operational user guidance shall identify all possible modes of operation of the TOE (including operation following failure or operational error), their consequences and implications for maintaining secure operation.

AGD_OPE.1.6c The operational user guidance shall, for each user role, describe the security measures to be followed in order to fulfill the security objectives for the operational environment as described in the ST.

AGD_OPE.1.7c The operational user guidance shall be clear and reasonable.

AGD_OPE.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.2.2 Preparative procedures (AGD_PRE.1)

AGD_PRE.1.1d The developer shall provide the TOE including its preparative procedures.

AGD_PRE.1.1c The preparative procedures shall describe all the steps necessary for secure acceptance of the delivered TOE in accordance with the developer's delivery procedures.

AGD_PRE.1.2c The preparative procedures shall describe all the steps necessary for secure installation of the TOE and for the secure preparation of the operational environment in accordance with the security objectives for the operational environment as described in the ST.
AGD_PRE.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
AGD_PRE.1.2e The evaluator shall apply the preparative procedures to confirm that the TOE can be prepared securely for operation.

5.2.3 Life-cycle Support (ALC)

5.2.3.1 Use of a CM system (ALC_CMC.2)

ALC_CMC.2.1d The developer shall provide the TOE and a reference for the TOE.
ALC_CMC.2.2d The developer shall provide the CM documentation.
ALC_CMC.2.3d The developer shall use a CM system.
ALC_CMC.2.1c The TOE shall be labelled with its unique reference.
ALC_CMC.2.2c The CM documentation shall describe the method used to uniquely identify the configuration items.
ALC_CMC.2.3c The CM system shall uniquely identify all configuration items.
ALC_CMC.2.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.3.2 Parts of the TOE CM coverage (ALC_CMS.4)

ALC_CMS.2.1d The developer shall provide a configuration list for the TOE.
ALC_CMS.2.1c The configuration list shall include the following: the TOE itself; the evaluation evidence required by the SARs; and the parts that comprise the TOE.
ALC_CMS.2.2c The configuration list shall uniquely identify the configuration items.
ALC_CMS.2.3c For each TSF relevant configuration item, the configuration list shall indicate the developer of the item.
ALC_CMS.2.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.3.3 Delivery procedures (ALC_DEL.1)

ALC_DEL.1.1d The developer shall document procedures for delivery of the TOE or parts of it to the consumer.
ALC_DEL.1.2d The developer shall use the delivery procedures.
ALC_DEL.1.1c The delivery documentation shall describe all procedures that are necessary to maintain security when distributing versions of the TOE to the consumer.
ALC_DEL.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.3.4 Flaw reporting procedures (ALC_FLR.2)

ALC_FLR.2.1d The developer shall document flaw remediation procedures addressed to TOE developers.
ALC_FLR.2.2d The developer shall establish a procedure for accepting and acting upon all reports of security flaws and requests for corrections to those flaws.
ALC_FLR.2.3d The developer shall provide flaw remediation guidance addressed to TOE users.
ALC_FLR.2.1c The flaw remediation procedures documentation shall describe the procedures used to track all reported security flaws in each release of the TOE.
ALC_FLR.2.2c The flaw remediation procedures shall require that a description of the nature and effect of each security flaw be provided, as well as the status of finding a correction to that flaw.
ALC_FLR.2.3c The flaw remediation procedures shall require that corrective actions be identified for each of the security flaws.
ALC_FLR.2.4c The flaw remediation procedures documentation shall describe the methods used to provide flaw information, corrections and guidance on corrective actions to TOE users.
ALC_FLR.2.5c The flaw remediation procedures shall describe a means by which the developer receives from TOE users reports and enquiries of suspected security flaws in the TOE.

ALC_FLR.2.6c The procedures for processing reported security flaws shall ensure that any reported flaws are remediated and the remediation procedures issued to TOE users.
ALC_FLR.2.7c The procedures for processing reported security flaws shall provide safeguards that any corrections to these security flaws do not introduce any new flaws.
ALC_FLR.2.8c The flaw remediation guidance shall describe a means by which TOE users report to the developer any suspected security flaws in the TOE.
ALC_FLR.2.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.4 Tests (ATE)

5.2.4.1 Evidence of coverage (ATE_COV.1)

ATE_COV.1.1d The developer shall provide evidence of the test coverage.
ATE_COV.1.1c The evidence of the test coverage shall show the correspondence between the tests in the test documentation and the TSFIs in the functional specification.
ATE_COV.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.4.2 Functional testing (ATE_FUN.1)

ATE_FUN.1.1d The developer shall test the TSF and document the results.
ATE_FUN.1.2d The developer shall provide test documentation.
ATE_FUN.1.1c The test documentation shall consist of test plans, expected test results and actual test results.
ATE_FUN.1.2c The test plans shall identify the tests to be performed and describe the scenarios for performing each test. These scenarios shall include any ordering dependencies on the results of other tests.
ATE_FUN.1.3c The expected test results shall show the anticipated outputs from a successful execution of the tests.
ATE_FUN.1.4c The actual test results shall be consistent with the expected test results.
ATE_FUN.1.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.4.3 Independent testing - sample (ATE_IND.2)

ATE_IND.2.1d The developer shall provide the TOE for testing.
ATE_IND.2.1c The TOE shall be suitable for testing.
ATE_IND.2.2c The developer shall provide an equivalent set of resources to those that were used in the developer's functional testing of the TSF.
ATE_IND.2.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
ATE_IND.2.2e The evaluator shall execute a sample of tests in the test documentation to verify the developer test results.
ATE_IND.2.3e The evaluator shall test a subset of the TSF interfaces to confirm that the TSF operates as specified.

5.2.5 Vulnerability assessment (AVA)

5.2.5.1 Vulnerability analysis (AVA_VAN.2)

AVA_VAN.2.1d The developer shall provide the TOE for testing.
AVA_VAN.2.1c The TOE shall be suitable for testing.
AVA_VAN.2.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
AVA_VAN.2.2e The evaluator shall perform a search of public domain sources to identify potential vulnerabilities in the TOE.

AVA_VAN.2.3e The evaluator shall perform an independent vulnerability analysis of the TOE using the guidance documentation, functional specification, TOE design and security architecture description to identify potential vulnerabilities in the TOE.

AVA_VAN.2.4e The evaluator shall conduct penetration testing, based on the identified potential vulnerabilities, to determine that the TOE is resistant to attacks performed by an attacker possessing Basic attack potential.

6 TOE Summary Specification

This chapter describes the security functions of the TOE.

6.1 TOE Security Functions

6.1.1 Security audit

Audit data is stored in memory and is separated into three types of logs; events, traffic logs, and self logs. Events are system-level notifications and alarms which are generated by the system to indicate events such as configuration changes, network attacks detected, or administrators logging in our out of the device. Traffic logs are directly driven by policies that allow traffic to go through the device. Self logs store information on traffic that is dropped and traffic that is sent to the device. Both audit events and traffic messages can be further defined depending on the severity of the message and/or event.

When logging and counting are enabled for a policy, all traffic will be logged to the traffic log.

There are two buffers for event logs, one for basic logs and one for alarms. There are also two buffers for traffic and self logs: one for traffic/self logs for traffic information and one for traffic/self events or alarms. The first tracks network traffic while the second stores information on alarms. Traffic/self alarms can be set in the policy such that when more traffic matches the policy than is configured in the policy alarm field, then an alarm will be logged.

The TOE can simultaneously send audit records to SDRAM and a remote syslog as a backup device to the audit log and an administrator controls this backup. The platform and storage device that control the syslog are not part of the TOE.

The TOE supports the following two configuration options. These two configuration options affect the overall handling of audit data. Refer to Table 6-1 for a summary of the possible TOE behaviors as controlled by these options.

Backup Syslog Device

The TOE can operate with only its internal memory used to store audit records, or it can be configured to send audit data to a backup syslog device. Security appliances provide memory to hold a fixed maximum number of audit records (many thousands of records on large models and about a thousand on smaller models). Memory is used because of the very high traffic flow speeds supported by a security appliance. Storing audit records on a disk or other permanent storage media simply is too slow to capture audited events and audit data would be lost using a slower audit recording device.

If a backup syslog device is not enabled (i.e., disabled), then once the storage limit is reached, the audit mechanism 'wraps' or acts as a first-in-first-out (FIFO) stack. The TOE overwrites the oldest audit information in the log memory buffer with the new audit information. When a backup syslog device is enabled, security appliances follow every write of an audit record to memory with an asynchronous write to a backup syslog device. This way memory acts as a high-speed FIFO buffer device to store megabytes of audit information, and the backup syslog device offers a longer term storage option.

Audit-Loss-Mitigation

Another configuration option supported by the TOE is called Audit-Loss-Mitigation. When Audit-Loss-Mitigation is disabled, the TOE overwrites existing audit records when the memory buffers used to store audit data wraps. Alternately, when Audit-Loss-Mitigation is enabled the TOE will stop processing traffic that may generate audit data when the memory buffers become full.

The technique of overwriting the oldest audit records once memory no longer has space for audit information limits the audit records that can be lost. All audit information is written at a speed that is directly proportional to audited activity. Audited activity on a protected network is rarely continuous over time, but occurs in bursts of high traffic and lulls where traffic that causes audited events are low. The worst case for audit loss would occur if the TOE

wrote an audit record in the last available memory location, and a burst of audited events occurred before the audit records in memory could be written to the backup syslog device. By overwriting the oldest audit information with the latest audit information to a very high-speed memory, the TOE does not completely lose any given audit record. That is, all audit records can be written to memory, and thus handled per TOE configuration. Additionally, the security appliances can be configured to notify the administrator (via a console alarm message) when the percentage of audit records that have not been written to the remote syslog has reached a specified percentage.

There is an internal field that identifies when an audit record has been written to the syslog device. When Audit-Loss-Mitigation is enabled, if this field indicates that the record has not been written to the syslog device, and the record is about to be overwritten, then an alarm will be created and all traffic will stop until all of the existing audit records are written to the syslog device. Once all existing audit records are written to the syslog device, network traffic is allowed to resume. During this stoppage of network traffic, device administration is allowed to continue, allowing an authenticated administrator to make configuration changes if necessary to prevent further problems with audit loss, such as changing an information flow policy. This feature ensures that no auditable events, except those taken by the authorized administrator will occur.

	Backup Syslog Device Enabled	**Backup Syslog Device Disabled**
Audit-Loss-Mitigation Disabled	This configuration operates the audit memory buffer as a FIFO stack, while sending audit data to an backup syslog device. Provided traffic can be sent to the backup syslog device as a sufficiently high speed this minimizes potential loss of audit data to conditions of extended bursts of audit generating traffic.	This configuration allows the audit memory buffer to operate as a FIFO stack and does not send audit records to an external server. This configuration presents the greatest potential for loss of audit records and prevents the audit data from creating an opportunity for denial of service.
Audit-Loss-Mitigation Enabled	This configuration operates the audit memory buffer as a FIFO stack, while sending audit data to a backup syslog device. However, when the buffer becomes full, traffic handling is stopped until all audit data in the audit memory buffer is sent to the backup syslog device.	This configuration prevents audit data loss when the audit memory buffer becomes full, but requires frequent monitoring and clearing of the audit trail by an administrator for continued TOE operation.

Table 6-1 Audit Storage Control Options

The Security, Audit, and Cryptographic administrators have access to the audit logs where the audit logs are stored while they are within the TOE. Only the Audit administrator may delete audit log entries. The available commands do not permit any user, including an authorized administrator to change the audit logs or permit restoration of the audit logs.

All log entries indicate the identity of the user who initiated the event, if applicable. Additionally, the information contained in the logs includes:

date:	The generation date of the event
time:	The generation time of the event
module:	The module name which generated the event
severity level:	The severity level of the event
type:	The event type
description:	The detailed description for the event may include the following:
ack-id:	The unique index of the security alarm event, from 1~4G.
user-id:	User id who generated the event
src-ip:	The source IP address which generated the event
dst-ip:	The destination IP address which generated the event
service port:	The service port which generated event

rule-id:	Rule id leading to this event
interface:	Interface of the event
outcome:	Success or failure of the event

The logs can contain the following auditable events:

1. Start-up and shutdown of the audit functions
2. Potential security violation was detected including the identification of what caused the generation of the alarm
3. Enabling and disabling of any of the analysis mechanisms
4. Opening the audit trail
5. Unsuccessful attempts to read information from the audit records
6. All modifications to the audit configuration that occur while the audit collection functions are operating
7. Actions taken due to exceeding the audit threshold
8. Actions taken due to the audit storage failure
9. Attempts at the generation and loading of a crypto key.
10. Generation and loading of key pair for digital signatures.
11. Failure of a cryptographic operation including type of cryptographic operation and applicable cryptographic mode of operation (no sensitive information is included in the audit record)
12. Changes to the pre-shared key used for authentication
13. All modifications to the key lifetimes.
14. Failure of the authentication in IKE Phase 1.
15. Failure to negotiate a security association in IKE Phase 2.
16. Decisions to permit or deny information flows.
17. Operation applied to each information flow permitted.
18. The reaching of the threshold for the unsuccessful authentication attempts
19. The actions (e.g. disabling of an account) taken
20. Identity of the unsuccessfully authenticated user and the identity of the administrator performing the function.
21. Successful and unsuccessful use of authentication mechanisms
22. All use of the local authentication mechanism
23. All use of the user identification mechanism used for authorized users (that is, those that authenticate to the TOE)
24. Success and failure of binding of user security attributes to a subject
25. All modifications in the behavior of the functions in the TSF
26. Enabling or disabling of the key-generation self-tests
27. All modifications in the behavior of the functions in the TSF
28. All manipulation of the security attributes
29. All modifications of the values of TSF data by the administrator
30. All modifications of the values of cryptographic security data by the cryptographic administrator
31. All modifications to the time and date used to form the time stamps by the administrator
32. All modifications to the information flow policy ruleset by the Security Administrator
33. All modifications of quota limits.
34. Actions taken when the quota is exceed (include the fact that the quota was exceeded)
35. All attempts to revoke security attributes
36. Modifications to the group of users that are part of a role
37. The fact that a failure or service discontinuity occurred
38. Resumption of the regular operation
39. Notification that a replay event occurred
40. Changes to the time
41. Execution of this set of TSF self tests
42. Locking of an interactive session by the session locking mechanism.
43. Successful unlocking of an interactive session.
44. Any attempts at unlocking an interactive session.
45. The termination of a remote session by the session locking mechanism
46. Denial of a session establishment due to the session establishment mechanism.

47. All attempts at establishment of a user session.
48. All attempted uses of the trusted channel functions.
49. Identifier of the initiator and target of all trusted channel functions.
50. All attempted uses of the trusted path functions.
51. Identification of the user associated with all trusted path invocations, if available.

The CLI provides an authorized administrator the ability to use 'set' commands to configure a security appliance, 'get' commands to display system configuration parameters and data, and 'clear' commands to remove data collected in various tables, memory, and buffers. The 'set' commands are used to set auditable events. The 'get log' command displays all records in the log. The 'get log' command can also be used to display records matching attributes of each audited event, including: user identity, source subject identity, destination subject identity, rule identity, ranges of dates, ranges of times, subject service identifiers, transport layer protocol or TOE network interfaces. The 'set log' command allows the security administrator to exclude certain events from being logged, based on the specific attributes of the audited event. Those attributes are:

1. user identity
2. event type
3. network identifier
4. subject service identifier
5. success of auditable security events
6. failure of auditable security events, and
7. rule identity.

Messages are reported by type and severity. For every log message within a message type, the message is documented, as well as the meaning of the message, and the appropriate action that an administrator needs to take. There are dozens of specific message types. 'Authentication' is but one type. Authentication message types relate to user authentication. Within this message there are four levels of severity: 1 - alert, 2 - warning, 3 - information, and 4 - notification.

The Security audit function is designed to satisfy the following security functional requirements:

- FAU_GEN.1: The TSF generates audit records for the events listed above. Audit records contain the identity of the subject that caused the event, date and time of the event, the event type, the success or failure of the event, and other event specific information that contains at least the information represented above.
- FAU_SAR.1: Administrators can read audit data using command line operations.
- FAU_SAR.3: The 'get log' command provided by the CLI provides the appropriate administrator the tools to review the audit logs, as well as to search and/or sort by attributes of each audited event.
- FAU_STG.1: The TOE provides the ability to delete audit log entries only to the Audit administrator. No other administrator may modify the audit trail because no interface is provided for such actions.
- FAU_STG.4: The TSF allows the security administrator to configure the system such that the TOE behaves in the following manner.
    - In the event that the TOE is configured to prevent the overwriting of the oldest records when local audit storage is consumed, the TSF prevents all traffic forwarding until audit storage is explicitly cleared by the audit administrator.

6.1.2 Cryptographic support

The TOE meets FIPS 140-2 requirements by allowing the administrator to enable a FIPS operating mode. The evaluated configuration of the TOE requires the use of this FIPS operating mode. The Cryptographic security function is described in the context of how it satisfies the cryptographic security requirements.

All security appliances comprising the TOE are FIPS 140-2 Security Level 2 validated. The certificate numbers for the completed FIPS evaluation are:

- SSG 5 and SSG 20, #<<Pending>>
- SSG 140, #<Pending>>
- SSG 320M and SSG 350M, #<Pending>>
- SSG 520M and SSG 550M, #<Pending>>
- Netcreen-ISG1000 and NetScreen-ISG2000, #<Pending>>
- NetScreen-5200 and NetScreen-5400, #<Pending>>

A VPN utilizes AES encryption and digital signature operations using teh FIPS-approved cryptomodule. The cryptomodule implements AES encryption using 128, 192 and 256 bit keys.

The cryptomodule also implements ECDSA using a base point of 256-bits or greater (as specified by the cryptographic administrator) for digital signature generation and verification. The FIPS-approved ECDSA algorithm is defined by ANSI X9.62-1998. The cryptomodule supports a single prime field (F sub p) named curve: secp256r1. The cryptomodule allows the administrator to generate an EC key-pair, create a certificate request for that key-pair, and configure a VPN using a certificate to digitally sign and verify communication. The cryptomodule does not support public key validation.

The encrypted communication channel between the TSF and an administrator in the environment is provided by the use of an administrator-initiated SSH session using public key certificate based authentication. This protocol provides encryption of the transmitted data that utilizes the FIPS cryptographic module for AES encryption using 128, 192 or 256 bit keys. ScreenOS ensures that administrators must use SSH to protect all remote administration sessions.

The Cryptographic support function is designed to satisfy the following security functional requirements:

- FCS_COP.1(1): The cryptomodule supports a FIPS-approved implementation of AES-CBC, using 128, 192 and 256 bit keys.
- FCS_COP.1(2): The cryptomodule supports a FIPS-approved implementation of AES-CBC, using 128, 192 and 256 bit keys. The cryptomodule also supports ECDSA with a key size of 256 bits using the NIST curve, P-256 and SHA256.

6.1.3 User data protection

Security appliances act as stateful inspection firewalls that examine each packet and track application-layer information for each connection by setting up a state table that spans multiple packets. This is used to determine whether incoming packets are legitimate. It eliminates the requirement to establish a TCP session with the firewall itself to access a service on the other side of the firewall (i.e. proxy the service).

By default, a security appliance denies all traffic in all directions.

6.1.3.1 VPN Policies

The VPN SFP by default enforces the use of an “access policy” that is established by an administrator to filter certain objects and to take an appropriate action depending upon the contents of a packet, or to apply a default policy that is available. Each access policy contains at least the following elements:

Addresses and/or Address Zones (source and destination)

Transport Layer (protocol)

Interface (i.e., physical network port)

Tunnel interface on which the traffic arrives and departs

Service (A service is considered a protocol assigned to a port or as data specific to a service such as FTP-GET)

The service can be filtered using the Application Layer Gateway[4] (ALG) software component of the TOE. ALG intercepts and analyzes specified traffic, allocates resources, and enforces dynamic policies defined to permit or deny traffic passing through the TOE. Through support of the ALG, the TOE provides the capability to filter DNS, RSH, FTP, and HTTP services, as well as granular HTTP component blocking. HTTP component blocking allows the administrator to selectively choose which HTTP components (e.g., .exe files, .zip files) are to be blocked by the TOE.

The addresses and/or address groups may be used to map a network or a group of networks to a security zone. T his allows the administrator to configure a policy that applies to a specific network or to a group of networks, rather than having to write multiple policies to perform a similar task for a group of networks.

The access policy can be configured to control information flow based on all combinations of these elements. Access policies only apply to TCP and UDP transport layer protocols. Access policies may be configured to permit, deny (drop silently), reject (drop with error sent to source), nat (perform address translation), or tunnel (permit with encryption or decryption) information matching the policy. The VPN SFP supports all of these actions. However, the tunnel action is required for an external IT entity to successfully invoke the tunnel interface and establish a VPN connection. VPN connections cause the encryption and decryption of information as it flows into and out of the TOE. The TOE also supports establishing multiple tunnels to a single tunnel interface.

By default, a security appliance denies all traffic in all directions. Security appliances are designed to prevent inappropriate information flows since all information that flows from one zone to another must pass through the security appliance.

Any time an information flow request is received by the TOE, the TOE performs a policy lookup to determine how the requesting information flow should be treated.

If the information flow request initiating a VPN tunnel arrives on an internal network, the information flow may be permitted to traverse through the TOE to another connected network if:

- the external IT entity initiating the information flow has successfully authenticated to the TOE using a key associated with the VPN connection;
- all the information security attribute values are unambiguously permitted by the information flow security policy rules, where such rules may be composed from all possible combinations of the values of the information flow security attributes, created by the authorized administrator;
- the presumed address of the source subject, in the information, translates to an internal network address;
- and the presumed address of the destination subject, in the information, translates to an address on the other connected network.

If the information flow request initiating a VPN tunnel arrives on the external network, the information flow may be permitted to traverse through the TOE to another connected network if:

- the external IT entity initiating the information flow has successfully authenticated to the TOE using a key associated with the VPN connection;
- all the information security attribute values are unambiguously permitted by the information flow security policy rules, where such rules may be composed from all possible combinations of the values of the information flow security attributes, created by the authorized administrator;
- the presumed address of the source subject, in the information, translates to an external network address;
- and the presumed address of the destination subject, in the information, translates to a valid address on the other connected network.

The TOE first checks to see if the source and destination zones are the same or different.

[4] The RSH ALG filtering is not supported when used with port address translation.

- If the source and destination zones are different, then the TOE performs a policy lookup in the interzone policy set list, or
- If no interzone policy is defined to permit the requested information flow, then the information flow is dropped by the default deny policy.

In addition to the set of policy checks an information flow request is subjected to, the TOE also checks information flow requests against IP spoofing, broadcast packets and loopback packets.

An information flow request is detected as IP spoofing if the request arrives on an external TOE interface and the presumed address of the source subject is an external IT entity on an internal network, or if the request arrives on an internal TOE interface and the presumed address of the source subject is an external IT entity on the external network.

An information flow request is detected as a broadcast packet if the request arrives on either an internal or external IPv4 TOE interface and the presumed address of the source subject is an external IT entity on a broadcast network.

An information flow request is detected as a loopback packet if the request arrives on either an internal or external TOE interface and the presumed address of the source subject is an external IT entity on the loopback network.

In addition to the actions identified, a policy may also be configured to perform Policy-Based Address Translation on information matching such a policy and may also be configured to block or reassemble fragmented packets pertaining to HTTP or FTP services.

Policy-Based Address Translation may be performed on either the presumed source IPv4 address of the information or on the presumed destination IPv4 address of the information.

Policy-Based Address Translation that is applied to the presumed source IPv4 address of the information may be configured to perform any of the following types of address translation:

- NAT-Src from a DIP Pool with PAT
- NAT-Src from a DIP Pool without PAT
- NAT-Src from a DIP Pool with Address Shifting
- NAT-Src from the Egress Interface IPv4 Address.

Policy-Based Address Translation that is applied to the presumed destination address of the information may be configured to perform any of the following types of address translation:

- NAT-Dst to a Single Iv4P Address with Port Mapping
- NAT-Dst to a Single IPv4 Address without Port Mapping
- NAT-Dst from an IPv4 Address Range to a Single IPv4 Address
- NAT-Dst between IPv4 Address Ranges.

6.1.3.2 Firewall Policies

The UNAUTHENTICATED SFP by default enforce the use of an “access policy” that is established by an administrator to filter on certain objects and to take an appropriate action depending upon the contents of a packet, or to apply a default policy that is available. Each access policy contains at least the following elements:

- Addresses and/or Address Zones (source and destination)
- Transport Layer (protocol)
- Interface (i.e., physical network port)
- Service (A service is considered a protocol assigned to a port or as data specific to a service such as FTP-GET).

The service can be filtered using the Application Layer Gateway[5] (ALG) software component of the TOE. ALG intercepts and analyzes specified traffic, allocates resources, and enforces dynamic policies defined to permit or deny traffic passing through the TOE. Through support of the ALG, the TOE provides the capability to filter DNS, RSH, FTP, and HTTP services, as well as granular HTTP component blocking. HTTP component blocking allows the administrator to selectively choose which HTTP components (e.g., .exe files, .zip files) are to be blocked by the TOE.

The addresses and/or address groups may be used to map a network or a group of networks to a security zone. This allows the administrator to configure a policy that applies to a specific network or to a group of networks, rather than having to write multiple policies to perform a similar task for a group of networks.

The access policy can be configured to control information flow based on all combinations of these elements. Access policies only apply to TCP and UDP transport layer protocols. Access policies may be configured to permit, deny (drop silently), reject (drop with error sent to source), NAT (perform address translation), or tunnel (permit with encryption or decryption) information matching the policy. A firewall configured in transparent mode only supports the actions to permit, deny or reject. A firewall configured in route mode only supports the actions to permit, deny, reject or NAT.

By default, a security appliance denies all traffic in all directions. Security appliances are designed to prevent inappropriate information flows since all information that flows from one zone to another must pass through the security appliance.

In addition to the actions identified, a policy may also be configured to perform Policy-Based Address Translation on information matching such a policy and may also be configured to block or reassemble fragmented packets pertaining to HTTP or FTP services.

Policy-Based Address Translation may be performed on either the presumed source IPv4 address of the information or on the presumed destination IPv4 address of the information.

Policy-Based Address Translation that is applied to the presumed source IPv4 address of the information may be configured to perform any of the following types of address translation:

- NAT-Src from a DIP Pool with PAT
- NAT-Src from a DIP Pool without PAT
- NAT-Src from a DIP Pool with Address Shifting
- NAT-Src from the Egress Interface IPv4 Address.

Policy-Based Address Translation that is applied to the presumed destination IPv4 address of the information may be configured to perform any of the following types of address translation:

- NAT-Dst to a Single IPv4 Address with Port Mapping
- NAT-Dst to a Single Iv4P Address without Port Mapping
- NAT-Dst from an IPv4 Address Range to a Single IPv4 Address
- NAT-Dst between IPv4 Address Ranges.

Any time an information flow request is received by the TOE, the TOE performs a policy lookup to determine how the requesting information flow should be treated.

If the information flow request arrives on an internal network, the information flow may be permitted to traverse through the TOE to another connected network if:

- all the information security attribute values are unambiguously permitted by the information flow security policy rules, where such rules may be composed from all possible combinations of the values of the information flow security attributes, created by the authorized administrator;
- the presumed address of the source subject, in the information, translates to an internal network address;

[5] The RSH ALG filtering is not supported when used with port address translation.

- and the presumed address of the destination subject, in the information, translates to an address on the other connected network.

If the information flow request arrives on the external network, the information flow may be permitted to traverse through the TOE to another connected network if:

- all the information security attribute values are unambiguously permitted by the information flow security policy rules, where such rules may be composed from all possible combinations of the values of the information flow security attributes, created by the authorized administrator;
- the presumed address of the source subject, in the information, translates to an external network address;
- and the presumed address of the destination subject, in the information, translates to a valid address on the other connected network.

The TOE first checks to see if the source and destination zones are the same or different.

- If the source and destination zones are different, then the TOE performs a policy lookup in the interzone policy set list, or
- If the source and destination zones are the same, then the TOE performs a policy lookup in the intrazone policy set list.

If the TOE performs the interzone or intrazone policy lookup and does not find a match, then the TOE checks the global policy set (route mode only) list for a match.

- If the TOE performs the interzone and global policy lookups and does not find a match, then the TOE applies the default deny policy to the packet.

In addition to the set of policy checks an information flow request is subjected to, the TOE also checks information flow requests against IP spoofing, broadcast packets and loopback packets.

Prior to applying the information policy ruleset, the TOE completely reassembles fragmented packets. Whenever a packet is received that is not associated with an allowed established session (e.g., the SYN flag is set without the ACK flag being set), the information flow policy ruleset is applied to the packet. Otherwise, the TSF associates a packet with an allowed established session using the stateful packet attributes.

The stateful packet attributes maintained by the TOE for connection-oriented protocols (e.g., TCP) include sequence number, acknowledgement number, and flags (i.e., SYN, ACK, RST, and FIN). For connectionless protocols, the TOE maintains as stateful packet attributes the source and destination network identifiers as well as the source and destination service identifiers.

6.1.3.2.1 Security Zone

A security zone is a segment of network space to which security measures are applied. Multiple security zones can be configured on a single security appliance by sectioning the network into segments to which various security policies may be applied to satisfy the needs of each segment. At a minimum, two security zones must be identified, basically to protect one area of the network from the other. Many security zones can also be established to bring finer granularity to a network security design, without deploying multiple security appliances to do so.

Each security appliance is also configured with a Global Zone. A Global Zone is a security zone without a security zone interface. The Global Zone serves as a storage area for mapped IP (MIP) and virtual IP (VIP) addresses. The predefined Global zone address “Any” applies to all MIPs, VIPs, and other user-defined addresses set in the Global zone. Because traffic going to these addresses is mapped to other addresses, the Global zone does not require an interface for traffic to flow through it.

6.1.3.2.1.1 Security Zone Interface

A security zone interface is an interface in which information can be sent to and from a security zone. Security zones support five types of security zone interfaces, which include physical interfaces, subinterfaces, aggregate interfaces, redundant interfaces, and virtual security interfaces. However, the evaluated configuration of the TOE may only utilize the physical interfaces, aggregate interfaces, and redundant interfaces.

6.1.3.2.1.1.1 Physical Interface

Each physical network port on the security appliance represents a physical interface, and the name of the interface is predefined. The name of a physical interface is composed of the media type, slot number (for some security appliances), and port number, for example, ethernet3/2 or ethernet2. A physical interface can bind to any security zone where it acts as a doorway through which traffic enters and exits the zone. Without a physical interface, no traffic can access the zone or leave it.

6.1.3.2.1.1.2 Aggregate Interface

The Juniper Networks NetScreen-5000 series supports aggregate interfaces. An aggregate interface is the accumulation of two or more physical interfaces, each of which shares the traffic load directed to the IP address of the aggregate interface equally among them. By using an aggregate interface, the amount of bandwidth available to a single IP address can be increased. Also, if one member of an aggregate interface fails, the other member or members can continue processing traffic, although with less bandwidth than previously available.

6.1.3.2.1.1.3 Redundant Interface

A redundant interface consists of binding two physical interfaces together to create one redundant interface, which you can then bind to a security zone. One of the two physical interfaces acts as the primary interface and handles all the traffic directed to the redundant interface. The other physical interface acts as the secondary interface and stands by in case the active interface experiences a failure. If that occurs, traffic to the redundant interface fails over to the secondary interface, which becomes the new primary interface. The use of redundant interfaces provides a first line of redundancy before escalating a failover to the device level.

6.1.3.2.2 Tunnel Zone

The concept of a tunnel zone is described only as background to for policy enforcement based upon zones.

A tunnel zone is a logical segment that hosts one or more tunnel interfaces. A tunnel zone is conceptually affiliated with a security zone in a “child-parent” relationship. The security zone acting as the “parent” provides the firewall protection to the encapsulated traffic. The tunnel zone provides packet encapsulation/decapsulation, and by supporting tunnel interfaces with IP addresses and net masks that can host mapped IP (MIP) addresses and dynamic IP (DIP) pools, can also provide policy-based NAT services. The security appliance uses the routing information for the carrier zone to direct traffic to the tunnel endpoint. The default tunnel zone is Untrust-Tun, and it is associated with the Untrust zone. Other tunnel zones can be created and bound to other security zones, with a maximum of one tunnel zone per carrier zone per virtual system. Virtual systems, however, are outside the scope of the evaluated configuration.

6.1.3.2.3 Function Zone

The function zone is a zone that performs a specific function. Functional zones support five types of zones, which include null zones, MGT zones, HA zones, self zones, and VLAN zones. However, the evaluated configuration of the TOE may only utilize the null zones and self zones. Each zone exists for a single purpose, as explained below.

6.1.3.2.3.1 Null Zone

This zone serves as temporary storage for any interfaces that are not bound to any other zone.

6.1.3.2.3.2 Self Zone

This zone hosts the interface for remote management connections. When connecting to the security appliance via HTTP, SSH, or Telnet, the self zone is used. Remote management is supported in the evaluated configuration of the TOE only via SSH (HTTP and Telnet management interfaces are disabled).

6.1.3.3 TOE Services

The TOE provides network services for SSH in support of remote administration. The TOE can also support response to ICMP echo requests (a.k.a., ping). Features such as administration of the TOE using a Telnet connection or SNMP are not included in the evaluated configuration.

Traffic destined for the TOE may arrive

1. directly at the destination TOE interface;
2. encapsulated and delivered to any TOE tunnel interface, then routed to the destination interface; or
3. delivered in plaintext to any TOE interface, then routed to the destination interface.

Traffic that is destined for the TOE is handled this way:

1. The packet is delivered to the destination TOE interface.
2. If management traffic has not been enabled on that interface, the packet is dropped.
3. If management traffic has been enabled, the type of management traffic enabled on that interface (e.g., SSH, Telnet, ping, SNMP)[6] is compared against the type of traffic enabled. If they do not match, the packet is dropped.
4. If a traffic type matches, it is delivered to the appropriate management daemon.

The firewall policy ruleset is not applied to this traffic if it is delivered directly to the destination TOE interface in plaintext. If it is delivered encapsulated or to another TOE interface, the rule set will be applied before forwarding it to the destination TOE interface.

Unauthenticated ICMP echo communications directed at the TOE are received and acknowledged per the configuration defined by the security administrator. Policy-Based Address Translation is applied to IPv4 interfaces configured for NAT.

Features such as administration of the TOE using a Telnet connection or SNMP are not included in the evaluated configuration.

The User data protection function is designed to satisfy the following security functional requirements:

- FDP__IFC.1(1): The UNAUTHENTICATED SFP applies to traffic to or from a network interface configured in Transparent mode, Route mode, or NAT mode that is not using a VPN tunnel.
- FDP__IFF.1(1): In Transparent and NAT modes, the UNAUTHENTICATED SFP enforces the use of an 'access policy' that is established by an administrator to filter on certain objects and to take an appropriate action depending upon the contents of a packet, or to apply a default policy that is available. The 'access policy' may be configured to permit or deny information matching the policy. Security appliances support multiple policies based upon a 'zone', using the rules defined in the Firewall policies' section above. The administrator can display the current configuration using the 'get policy' command, then configure a single new policy on the command line, which they can view in context of the whole configuration before entering the command for the new policy.

  In Transparent and NAT modes, the actions defined in an 'access policy' that enforces the UNAUTHENTICATED SFP are mapped to actions in the corresponding SFRs as follows:

  - the pass action maps to permit; and
  - the drop action maps to both deny and reject.

  In Route mode, the UNAUTHENTICATED SFP enforces the use of an 'access policy' that is established by an administrator to filter on certain objects and to take an appropriate action depending upon the contents of a packet, or to apply a default policy that is available. The 'access policy' may be configured to permit, deny or NAT information matching the policy. Security appliances support multiple policies based upon a

[6] Only ping and SSH are services supported in an evaluated configuration. Telnet, HTTP and SNMP based management are excluded as management interfaces in an evaluated configuration.

'zone', using the rules defined in the Firewall policies' section above. The administrator can display the current configuration using the 'get policy' command, then configure a single new policy on the command line, which they can view in context of the whole configuration before entering the command for the new policy.

In Route mode, the actions defined in an 'access policy' that enforces the UNAUTHENTICATED SFP are mapped to actions in the corresponding SFRs as follows:

- the pass action maps to permit; and
- the drop action maps to both deny and reject.

Policy-Based Address Translation is applied to IPv4 interfaces configured for NAT.

- FDP_IFC.1(2): The VPN SFP applies to traffic to or from a network interface configured in Transparent mode, Route mode or NAT mode that is using a VPN tunnel.
- FDP_IFF.1(2): The VPN SFP is enforced on information flows matching an 'access policy' defined by an administrator. The 'access policy' may be configured to pass, drop or tunnel information matching the policy. Security appliances support multiple policies based upon a 'zone', using the rules and algorithm defined in the 'VPN Policies' section above. The administrator can display the current configuration using the 'get policy' command, then configure a single new policy on the command line, which they can view in context of the whole configuration before entering the command for the new policy. The actions defined in an 'access policy' that enforces the VPN SFP are mapped to actions in the corresponding SFRs as follows:
  - the pass action maps to permit;
  - the send IPSEC encrypted action maps to tunnel;
  - the decrypt, verify authentication and pass action maps to tunnel; and
  - the drop action maps to both deny and reject.
- FDP_RIP.2: There are only two resources made available to information flowing through a security appliance. One is the temporary storage of packet information when access is requested and when information is being routed. The second type of information is key material.

  To secure all connection attempts, security appliances use a dynamic packet filtering method known as stateful inspection. Using this method, a security appliance notes various components in a TCP packet header. State information recognized by the device includes: source and destination IP addresses, source and destination port numbers, packet sequence numbers, and packet length. The security appliance maintains the state of each TCP session traversing the firewall. This means that security appliances keep track of packet length and packet attributes such that each packet must be complete and correct for information to flow from source to destination. The security appliance interprets every byte in a complete information stream from the first packet to the last. All temporary storage is accounted for in that the size of a temporary storage relative to every packet is known. Therefore, no residual information from packets not associated with a specific information stream can traverse through a security appliance. =

6.1.4 Identification and authentication

The identification and authentication security function is described in the context of how it satisfies the identification and authentication security requirements.

The security appliance provides only a local authentication mechanism using a local user database for authentication. ScreenOS checks for an administrator account on the local database. The TSF enforces the authentication decision, allowing access to the system only when a positive authentication occurs.

All administrative actions require identification and authentication prior to the action being performed. Thus, the TSF can associate a user id with every administrative action. All such actions are tracked by logging the user id along with the actions they perform on the TOE. The user id is associated with a particular type of administrator role. Source IP addresses are associated with every packet received from an external IT entity. The administrator's

role and identity are explicitly established based on the user's security attributes at the time of login and cannot be changed within a session.

The Identification and authentication function is designed to satisfy the following security functional requirements:

- FIA_AFL.1: The Security Administrator has the ability to specify the number of unsuccessful login attempts allowed before the device closes a login session. The default is 3 and the range of allowed values is 1 through 255. This is only applicable to the remote administrators trying to authenticate to the security appliance remotely. The TSF has the ability to enforce the Security Administrator specified action when the maximum number of unsuccessful login attempts is reached or exceeded. The Security Administrator can either choose to prevent remote administrators from logging in for a specified time period or can totally disable the login until further action is taken by the security administrator.

  The TOE does not support authentication attempts by external entities such as an NTP server. All trusted channels/trusted paths to authorized servers must be initiated by the TOE and must protect TOE to server communications.

- FIA_ATD.1: The TSF maintains an identity and password for each administrator authorized to administer the security appliance. All non-root users fall into one of these three administrator categories (Security, Crypto and Audit). A new account is given root privileges by default, and must be explicitly assigned to one of the three administrative categories.

- FIA_UAU.1: Security appliances require administrative personnel to perform identification and authentication before they may access any of the TOE functions or data. Once their identity has been provided, the administrator must enter the correct password in order to be successfully authenticated. With respect to network servers that would obtain services from the TOE, only ICMP echo and ARP communications are unmediated and hence are unauthenticated. Network traffic flowing in accordance with the UNAUTHENTICATED SFP is not subject to authentication.

- FIA_UID.2: Users of the TOE include administrators and other network servers that would obtain services from the TOE. Administrators must present a valid user id during the login process. Network servers, and all other network entities including those sending traffic subject to the UNAUTHENTICATED SFP, are identified by their source IP address, which is provided as part of every IP packet.

- FIA_UAU.4: The encrypted communication path between the TSF and a remote administrator is provided by the use of a SSH session. Remote administrators of the TSF initiate communication with the TSF through the SSH tunnel. While the TOE also supports SSH password identification. The SSH protocol ensures that the data transmitted over a SSH session cannot be disclosed or altered by negotiating a session key for each new SSH connection that ensures that authentication data traversing the network (from remote administrator SSH client to the TOE) is uniquely encrypted during transmission, thus ensuring that authentication material that may be gathered by a network intermediary is used only a single time.

6.1.5 Security management

The factory default configuration of the TOE has a single administrative account ('root') with all privileges on the device. To place the TOE in the evaluated configuration, the administrator uses this root account to configure three accounts, corresponding to the Security, Cryptographic and Audit administrative roles. The TOE also recognizes non-administrative users who may require authentication prior to permitting their traffic to traverse the firewall. The only other form of entity recognized by the TOE is an external IT entity. An external IT entity may be either authorized or not authorized. Authorized IT entities are those external IT entities for which the TOE has been configured to utilize the external functionality (e.g., an external NTP server can be an authorized IT entity).

When each new account is created, an attribute must be assigned to that account indicating the role associated with that account. An attribute exists for each of the Root, Security, Cryptographic and Audit administrative roles. Once the an account has been created, there is no overlap between the privileges available to the security, cryptographic and Audit roles, except the ability to review the audit trail, start/stop the TOE, and invoke self-tests. Any account with "root" privileges is able to perform all operations.

Administrators in any of these roles can login to the TOE either locally or remotely. Use of the web interface is not included in the evaluated configuration. Remote administration utilizes an SSH protected communication pathway to present a command line interface.

The root administrator is allowed to perform all functions. The following are activities that only a root administrator can perform:

1. create, delete, modify, and view user security attributes

The security administrator is allowed to perform the following:

2. Start-up and shutdown of the TOE
3. enable, disable, determine and modify the set of audited events collected
4. perform searches and sorting of audit data
5. create, delete, modify and view the security attributes referenced in the UNAUTHENTICATED SFP policies
6. query, modify, delete, clear all TSF data, except cryptographic security data and audit data
7. specify alternative initial values to override the default values for the firewall UNAUTHENTICATED SFP while in transparent and route mode
8. enable or disable the use of the SSH mechanism to support remote administration
9. specify the addresses from which remote administration can occur
10. modify and set the threshold for the number of permitted authentication attempt failures from users;
11. restore authentication capabilities for users that have met or exceeded the threshold for permitted authentication attempt failures
12. enable and disable external IT entities from communicating to the TOE
13. backup of user attribute values, information flow security policy rules, and audit trail data
14. recover to the state following the last backup
15. change or set the time and date used to form the time stamps
16. execute TSF self tests.

The cryptographic administrator is allowed to perform the following:

1. Start-up and shutdown of the TOE
2. perform searches and sorting of audit data;
3. modify cryptographic security data; and
4. execute TSF self tests.

The audit administrator is allowed to perform the following:

1. Start-up and shutdown of the TOE
2. perform searches and sorting of audit data; and
3. delete audit log entries.
4. execute TSF self tests..

The Security management function is designed to satisfy the following security functional requirements:

- FMT_MOF.1: Only the security administrator, cryptographic administrator and audit administrator can perform the management operations as described above.
- FMT_MSA.3: By default, a security appliance denies all traffic in all directions. Only the Security administrator can specify policy rules that permit traffic to flow through the TOE.
- FMT_SMF.1: The TSF provides all of the management operations specified by the FMT_SMF.1 requirement as shown in the lists above.
- FMT_SMR.1: The TSF defines the roles Cryptographic, Audit and Security Administrator with duties as described above. Users in these roles may login to the TOE remotely. Each of these administrative roles is considered authorized, and together they provide the necessary management operations.

6.1.6 Protection of the TSF

For networks connected to the security appliance, all network traffic is routed through the security appliance. Once network traffic is received on one of the security appliance network ports, it is always subject to the security policy rules. The Protection of the TSF security function is described in the context of how it satisfies the Protection of the TSF security requirements.

Protection of the TOE from physical tampering is ensured by its environment. It is assumed that security appliances will remain attached to the physical connections made by an administrator so that an appliance cannot be bypassed. Each security appliance is completely self-contained. The hardware and firmware provided by security appliances provide all the services necessary to implement the TOE. There are no external interfaces into the TOE other than the physical ports provided. No general purpose operating system, disk storage, or programming interface is provided.

The TOE protects its management functions by isolating them through authentication. Any interface that is controlled by a security zone can have two IP addresses. One is a physical port interface IP address (or a logical sub-interface), which connects to a network. The other is a second logical IP address for receiving administrative traffic.

Administrators are instructed to change the default password. If an administrator forgets their password, the security appliance has to be reset to the factory settings and connection configurations and Access Policy profiles are lost.

Logically, each security appliance is protected by the integrity of the protocol interpreters supporting the external interface. As long as network packets remain objects to be operated on by ScreenOS, the TSF is protected. ScreenOS is a custom operating system that runs in hardware and firmware, remains memory resident, and supports only trusted processes. A security appliance provides no file abstractions or permanent storage for 'executables' to remain for further execution. ScreenOS has been designed to control the protocols that it recognizes at its external interfaces.

Each identification and authentication interface of the security appliance that provides access to TSF internal objects is password protected, physically protected, and only can be manipulated by a person acting in an administrative role.

The underlying operating system is a monolithic real time operating system that is purpose built and hence there are no untrusted processes running in the kernel memory that would tamper with other processes at any point in time.

The operating system by design provides memory protection by using virtual memory and paging. Each process normally runs in its own virtual memory space (either static (1-1) mapping or dynamic mapping), and, unless explicitly requested, cannot access the memory of other processes. This is the basis for memory protection in ScreenOS. The cryptographic processes also are protected by the same mechanism.

A timestamp is stored internally as a count of the number of clock ticks since the device booted. This value is guaranteed to be monotonically increasing. Timestamps are converted to calendar time for display purposes or when transmitting values externally, such as to syslog servers.

The Protection of the TSF function is designed to satisfy the following security functional requirements:

- FPT_STM.1: Security appliance hardware provides a reliable clock, and the ScreenOS uses this clock to provide reliable time stamps. Both are part of the TSF.

7 Protection Profile Claims

The TOE conforms to the following protection profile:

- U.S. Government Traffic-Filter Firewall Protection Profile For Basic Robustness Environments, Version 1.1, July 25, 2007 (TFFW BR PP).

7.1 Security Problem Definition

All of the security problem definition and objective statements have been drawn from a validated PP (the TFFW BR PP). Please consult that PP for the applicable correspondence rationale. This ST includes all of the Threats, Policies and Assumptions from the PP.

7.2 Security Functional Requirements

The following table summarizes how the security functional requirements (SFRs) from this security target (ST) correspond to the requirements from the TFFW BR PP.

ST SFR	Relation to TFFW BR PP SFR
FAU_GEN.1: Audit data generation	Requirement in PP is the same as this ST. An editorial change was made to re-numbering of the table containing audit records and to correct the corresponding references within the requirement to that table.
FAU_SAR.1: Audit review	Requirement in PP is the same as this ST
FAU_SAR.3: Selectable audit review	Requirement in PP is the same as this ST
FAU_STG.1: Protected audit trail storage	The Requirement in the ST has been changed to reflect the CC v3.1 wording which better states the meaning of the requirement than the wording used in the PP.
FAU_STG.4: Prevention of audit data loss	Requirement in PP is the same as this ST
FCS_COP.1(1): Cryptographic Operation (remote sessions)	Requirement in PP is the same as this ST
FCS_COP.1(2): Cryptographic Operation (VPN operations)	Requirement in this ST is in addition to functionality required by the PP.
FDP_IFC.1(1): Subset information flow control (Unauthenticated Policy)	Requirement in PP is the same as this ST. Operations left to the ST writer were completed and denoted as assignments.
FDP_IFF.1(1): Simple security attributes (Unauthenticated Policy)	Requirement in PP is the same as this ST.
FDP_IFC.1(2): Subset information flow control (VPN Policy)	Requirement in this ST is in addition to functionality required by the PP.
FDP_IFF.1(2): Simple security attributes (VPN Policy)	Requirement in this ST is in addition to functionality required by the PP.
FDP_RIP.2: Full residual information protection	This ST replaces FDP_RIP.1 from the PP with FDP_RIP.2 because it is hierarchically more comprehensive.
FIA_AFL.1: Authentication failure handling	This ST completes the operation in FIA_AFL.1.1 by allowing the administrator rather than the ST author to determine the “settable non-zero number” that specifies the number of failed authentication attempts to detect.

ST SFR	Relation to TFFW BR PP SFR
FIA_ATD.1: User attribute definition	This ST completes the operation in FIA_ATD.1.1 by including a password as authentication material retained for users.
FIA_UAU.1: Timing of authentication	Requirement in PP is the same as this ST
FIA_UID.2: User identification before any action	Requirement in PP is the same as this ST
FIA_UAU.4: Single-use authentication mechanisms	Requirement in PP is the same as this ST
FMT_MOF.1: Management of security functions behavior	Requirement in PP is the same as this ST. Operation left to the ST writer was completed and denoted as an assignment.
FMT_MSA.3: Static attribute initialization	Requirement in PP is the same as this ST
FMT_SMR.1: Security Roles	Requirement in PP is the same as this ST
FMT_SMF.1: Specification of Management Functions	Requirement is added to meet NIAP policy.
FPT_STM.1: Reliable Time Stamps	Requirement in PP is the same as this ST Note that the PP erroneously identifies this as FPT_RVM.1 in one place.

Table 7-1 Correspondence Rationale for ST and PP SFRs

7.3 Assurance Requirements

The assurance requirements in the TFFW BR PP are identical to those found in this Security Target.

8 Rationale

This section provides the rationale for completeness and consistency of the Security Target. The rationale addresses the following areas:

- Security Objectives;
- Security Functional Requirements;
- Requirement Dependencies;
- Extended Requirements;
- TOE Summary Specification; and,
- PP Claims.

8.1 Security Objectives Rationale

All of the security environment and objective statements except those listed below have been drawn from a validated PP (TFFW BR PP). Please consult that PP for the applicable correspondence rationale.

- O.INTEGRITY maps to the policy P.INTEGRITY. O.INTEGRITY satifies this policy by ensuring that all IPSEC encrypted data received from a peer TOE is properly decrypted and authentication verified.

8.2 Security Requirements Rationale

All of the security functional requirements except those listed below have been drawn from a validated PP (TFFW BR PP). Please consult that PP for the applicable rationale.

- The requirements FDP_IFC.1(2), FDP_IFF.1(2), and FCS_COP.1(2) map to O.INTEGRITY. These requirements satify this objective by ensuring that all IPSEC encrypted data received from a peer TOE is properly decrypted and authentication verified.

8.3 Requirement Dependency Rationale

All of the requirements have been drawn from a validated PP (TFFW BR PP). Please consult that PP for the applicable rationale. The following table includes only those SFRs that have been added to this Security Target.

ST Requirement	CC Dependencies	ST Dependencies
FCS_COP.1(2)	(FDP_ITC.1, FDP_ITC.2, or FCS_CKM.1) and FCS_CKM.4	See Note Below
FMT_SMF.1	None	None
FDP_IFC.1(2)	FDP_IFF.1	FDP_IFF.1(2)
FDP_IFF.1(2)	FDP_IFC.1 FMT_MSA.3	FDP_IFC.1(2) FMT_MSA.3

NOTE:

Section 6.5 of the "U.S. Government Protection Profile for Traffic Filter Fierewall In Basic Robustness Environments", contains the following justification for not satisfying all of the dependencies for FCS_COP.1.

> "Functional component FCS_COP.1(2) depends on the following functional components: FCS_CKM.1 Cryptographic key generation, and FCS_CKM.4 Cryptographic key destruction. The cryptographic modules of the TOE are FIPS PUB 140-2 compliant. Because the cryptographic module is indeed compliant with this FIPS PUB, then the dependencies of key generation, key destruction and secure key values will have been satisfied in becoming FIPS PUB 140-2 compliant. For more information, refer to section 4.7 of FIPS PUB 140-2. "

This same rationale is applied to the iteration of FCS_COP.1(2) which specifies the VPN encryption capabilities of the TOE.

8.4 Extended Requirements Rationale

There are no extended requirements in this Security Target.

8.5 TOE Summary Specification Rationale

Each subsection in Section 6, the TOE Summary Specification, describes a security function of the TOE. Each description is followed with rationale that indicates which requirements are satisfied by aspects of the corresponding security function. The set of security functions work together to satisfy all of the security functions and assurance requirements. Furthermore, all of the security functions are necessary in order for the TSF to provide the required security functionality.

This Section in conjunction with Section 6, the TOE Summary Specification, provides evidence that the security functions are suitable to meet the TOE security requirements. The collection of security functions works together to provide all of the security requirements. The security functions described in the TOE summary specification are all necessary for the required security functionality in the TSF. **Table 8-1 Security Functions vs. Requirements Mapping** demonstrates the relationship between security requirements and security functions.

	Security audit	Cryptographic Support	User data protection	Identification and authentication	Security management	Protection of the TSF
FAU_GEN.1	X					
FAU_SAR.1	X					
FAU_SAR.3	X					
FAU_STG.1	X					
FAU_STG.4	X					
FCS_COP.1(1)		X				
FCS_COP.1(2)		X				
FDP_IFC.1(1)			X			
FDP_IFF.1(1)			X			
FDP__IFC.1(2)			X			
FDP__IFF.1(2)			X			
FDP_RIP.2			X			
FIA_AFL.1				X		
FIA_ATD.1				X		
FIA_UAU.1				X		
FIA_UID.2				X		
FIA_UAU.4				X		
FMT_MOF.1					X	
FMT_MSA.3					X	

	Security audit	Cryptographic Support	User data protection	Identification and authentication	Security management	Protection of the TSF
FMT_SMR.1					X	
FMT_SMF.1					X	
FPT_STM.1						X

Table 8-1 Security Functions vs. Requirements Mapping

8.6 PP Claims Rationale

See Section 7, Protection Profile Claims.