

Juniper Networks Security Appliances Security Target

Version 2.0
March 5, 2010

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

1. SECURITY TARGET INTRODUCTION	5
1.1 SECURITY TARGET, TOE AND CC IDENTIFICATION	5
1.2 CONFORMANCE CLAIMS	6
1.3 CONVENTIONS	6
1.4 TERMINOLOGY	7
2. TOE DESCRIPTION	12
2.1 TOE OVERVIEW	12
2.2 PRODUCT DESCRIPTION	12
2.2.1 Hardware	13
2.2.2 ScreenOS	13
2.2.3 Policies	13
2.2.4 Services	14
2.3 TOE CONFIGURATIONS	14
2.3.1 Interface Modes	14
2.3.2 VPN	17
2.4 TOE ARCHITECTURE	18
2.4.1 Physical Boundaries	18
2.4.2 Logical Boundaries	18
2.5 TOE DOCUMENTATION	22
3. SECURITY ENVIRONMENT	23
3.1 ORGANIZATIONAL POLICIES	23
3.2 THREATS	23
3.3 ASSUMPTIONS	24
4. SECURITY OBJECTIVES	25
4.1 SECURITY OBJECTIVES FOR THE TOE	25
4.2 SECURITY OBJECTIVES FOR THE ENVIRONMENT	26
5. IT SECURITY REQUIREMENTS	28
5.1 TOE SECURITY FUNCTIONAL REQUIREMENTS	28
5.1.1 Security audit (FAU)	31
5.1.2 Cryptographic support (FCS)	33
5.1.3 User data protection (FDP)	37
5.1.4 Identification and authentication (FIA)	44
5.1.5 Security management (FMT)	45
5.1.6 Protection of the TSF (FPT)	49
5.1.7 Resource utilization (FRU)	50
5.1.8 TOE access (FTA)	50
5.1.9 Trusted path/channels (FTP)	51
5.2 TOE SECURITY ASSURANCE REQUIREMENTS	52
5.2.1 Development (ADV)	52
5.2.2 Guidance documents (AGD)	54
5.2.3 Life-cycle Support (ALC)	55
5.2.4 Tests (ATE)	57
5.2.5 Vulnerability assessment (AVA)	57
6. TOE SUMMARY SPECIFICATION	59
6.1 TOE SECURITY FUNCTIONS	59
6.1.1 Security audit	59
6.1.2 Cryptographic support	63
6.1.3 User data protection	66
6.1.4 Identification and authentication	72

6.1.5	<i>Security management</i>	73
6.1.6	<i>Protection of the TSF</i>	75
6.1.7	<i>Resource utilization</i>	76
6.1.8	<i>TOE access</i>	77
6.1.9	<i>Trusted path/channels</i>	77
7.	PROTECTION PROFILE CLAIMS	79
7.1	SECURITY ENVIRONMENT.....	79
7.2	SECURITY FUNCTIONAL REQUIREMENTS.....	81
7.3	ASSURANCE REQUIREMENTS.....	85
8.	RATIONALE	86
8.1	SECURITY OBJECTIVES RATIONALE.....	86
8.2	SECURITY REQUIREMENTS RATIONALE.....	86
8.3	REQUIREMENT DEPENDENCY RATIONALE.....	86
8.4	EXTENDED REQUIREMENTS RATIONALE	86
8.5	TOE SUMMARY SPECIFICATION RATIONALE.....	86
8.6	PP CLAIMS RATIONALE.....	88
9.	AUDIT EVENTS	89
10.	STATISTICAL RANDOM NUMBER GENERATOR TESTS	96

LIST OF TABLES

Table 1	TOE Security Functional Requirements	31
Table 2	Medium Robustness Assurance Requirements	52
Table 3	PP Correspondence Rationale for Threats	79
Table 4	PP Correspondence Rationale for Security Policies	80
Table 5	PP Correspondence Rationale for Assumptions	80
Table 6	PP Correspondence Rationale for Security Objectives	80
Table 7	PP Correspondence Rationale for Objectives for the Environment	81
Table 8	PP Correspondence Rationale for SFRs	81
Table 9	Security Functions vs. Requirements Mapping	88

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 Security Appliances provided by Juniper Networks. The security appliances Target of Evaluation (TOE) primarily supports the definition of and enforces information flow policies among network nodes. The security appliance provides for stateful inspection of every packet that traverses the network via the TOE. All information flow from one network node to another passes through a security appliance, 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 its security functions.

The Security Target contains the following additional sections:

- TOE Description (Section 2)
- Security Environment (Section 3)
- Security Objectives (Section 4)
- IT Security Requirements (Section 5)
- TOE Summary Specification (Section 6)
- Protection Profile Claims (Section 7)
- Rationale (Section 8)
- Audit Events (Section 9)
- Statistical Random Number Generator Tests (Section 10)

1.1 Security Target, TOE and CC Identification

ST Title – Juniper Networks Security Appliances Security Target

ST Version – Version 2.0

ST Date – March 5, 2010

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	6.2.0r3a
Juniper Networks NetScreen ISG 2000	NS-ISG-2000, NS-ISG-2000-DC, NS-ISG-2000B, NS-ISG-2000B-DC	6.2.0r3a
Juniper Networks NetScreen 5200	NS-5200, NS-5200-DC	6.2.0r3a
Juniper Networks NetScreen 5400	NS-5400, NS-5400-DC	6.2.0r3a
Juniper Networks SSG5 Secure Services Gateway	SSG-5-SB, SSG-5-SH	6.2.0r3

Juniper Networks SSG20 Secure Services Gateway	SSG-20-SB, SSG-20-SH	6.2.0r3
Juniper Networks SSG140 Secure Services Gateway	SSG-140-SB, SSG-140-SH	6.2.0r3
Juniper Networks SSG320M Secure Services Gateway	SSG-320M-SH, SSG-320M-SH-N-TAA, SSG-320M-SH-DC-N-TAA	6.2.0r3
Juniper Networks SSG350M Secure Services Gateway	SSG-350M-SH, SSG-350M-SH-N-TAA, SSG-350M-SH-DC-N-TAA	6.2.0r3
Juniper Networks SSG520M Secure Services Gateway	SSG-520M-SH, SSG-520M-SH-N-TAA, SSG-520M-SH-DC-N-TAA	6.2.0r3
Juniper Networks SSG550M Secure Services Gateway	SSG-550M-SH, SSG-550M-SH-N-TAA, SSG-550M-SH-DC-N-TAA	6.2.0r3

TOE Developer – Juniper Networks

Evaluation Sponsor – Juniper Networks

CC Identification – Common Criteria for Information Technology Security Evaluation, Version 3.1, September, 2007

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, September, 2007.
 - Part 2 Extended
- Common Criteria for Information Technology Security Evaluation Part 3: Security Assurance Requirements, Version 3.1, September, 2007.
 - Part 3 Conformant - EAL 4
 - Assurance Level: EAL4 augmented with ADV_FSP.5, ADV_INT.3, ADV_TDS.4, ALC_FLR.2, ATE_DPT.3, and AVA_VAN.3

The TOE meets all of the security requirements of the following Protection Profiles, except for AVA_CCA_(EXT).1 and AVA_VAN.4:

- U.S. Government Virtual Private Network (VPN) Boundary Gateway Protection Profile For Medium Robustness Environments, Version 1.2, 30 January 2009 (VPN PP).¹
- U.S. Government Traffic-Filter Firewall Protection Profile For Medium Robustness Environments, Version 1.1, July 25, 2007 (TFFW PP).

1.3 Conventions

Since this security target is claiming compliance with two PPs, 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 PPs 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)

¹ The testing associated with AVA_CCA_(EXT).1 and AVA_VAN.4 as required by these Protection Profiles is pending completion by NSA.

and formatting (e.g., special fonts) used within the PPs to designate operations are being removed by 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 Medium Robustness VPN 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.

In many cases, a requirement in the TFFW PP is identical to a requirement in the VPN PP. When these two PPs change a requirement in a significantly different manner, this security target will restate the requirement such that all aspects of both versions of the requirement are specified, thus making the version of the requirement in this ST a superset of the required functionality from these two PPs. For requirements that are in the TFFW PP and not in the VPN PP this security target will append “_(FW)” to the requirement name. The exception to this is the FDP_IFC and FDP_IFF.1 requirements. Simple iterations have been performed upon these requirements because the SFP defined within these requirements makes it clear which iteration is from the VPN PP and which is from the TFFW PP. For all of the TFFW specific requirements (i.e., those denoted with “_(FW)”), the assignment, selection and refinement conventions mentioned above will be used to denote operations performed by this ST with respect to the requirement as found in the TFFW PP.

The CC paradigm also allows protection profile and security target authors to create their own requirements. Such requirements are termed ‘extended requirements’ and are permitted if the CC does not offer suitable requirements to meet the authors’ needs. Extended requirements must be identified and are required to use the CC class/family/component model in articulating the requirements. In these PPs, extended requirements are indicated with the “(EXT)” following the component name. This security target reproduces that convention.

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. 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. 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.
Destination Network Address Translation (NAT-dst)	
Dynamic IP (DIP) Pool	
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. 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.
FIPS 140-2	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.
Firmware	
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.
IP Security (IPSEC)	An IP security protocol that provides for encapsulation of standard IP packets into Type 51 IP, allowing firewalls to recognize and admit encapsulated, encrypted data.

Mapped IP Address (MIP)	<p>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.</p>
Network Address Translation (NAT)	<p>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.</p>
NAT Source (NAT-src)	<p>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.</p>
Network Basic Input/Output System (NetBIOS)	<p>An application programming interface used in conjunction with other programs to transmit messages between applications running on PCs hooked to a local area network.</p>
Network	<p>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.</p>
Node	<p>A concentration point in a network where numerous trunks come together at the same switch.</p>
Packet	<p>A block of data sent over the network transmitting the identities of the sending and receiving stations, error-control information, and message.</p>
Port Address Translation (PAT)	<p>The translation of the original source port number in a packet to a different, randomly designated port number.</p>
Public-Key Infrastructure (PKI)	<p>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.</p>
Session	<p>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.</p>
Session Table	<p>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.</p>

Stateful inspection	<p>Also referred to as <i>dynamic packet filtering</i>. 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.</p>
Synchronous Dynamic Random Access Memory (SDRAM)	<p>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).</p>
Tampering	<p>An unauthorized modification that alters the proper functioning of equipment or a system in a manner that degrades the security or functionality it provides.</p>
Transmission Control Protocol/Internet Protocol (TCP/IP)	<p>A communications protocol developed under contract from the U.S. Department of Defense to interconnect 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.</p>
Tunneling	<p>Use of one data transfer method to carry data for another method.</p>
User Datagram Protocol (UDP)	<p>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.</p>
Virtual IP Address (VIP)	<p>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.</p>
Virtual Private Network (VPN)	<p>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.</p>

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 Security Appliances.

Juniper Networks Security Appliances, hereafter referred to as security appliances, are integrated security network 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 consist of integrated security network appliances that operate as a central security hub in a networked configuration. The security appliances control traffic flow through the network. 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. All security appliances have hardware accelerated IPsec encryption and very low latency, allowing them to fit into any network. Installing and managing appliances is accomplished using a command line interface (CLI).

The TOE includes the security appliances that run ScreenOS, a custom operating system. The security appliances that meet the definition of TOE include several models. Each identified model consists of hardware and ScreenOS that runs in firmware. The set of models included are identified in Section 1.1 of this document.

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 relies on a virtual router (VR) component, which functions as a router and has its own interfaces and its own routing table. In ScreenOS, a security appliance 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 to be 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, 3DES and SHA-1. The remaining cryptographic and firewall functionality is performed in software. The SSG 140, 320M, 350M, 520M and 550M use the Cavium Nitrox Lite ASIC
- The SSG 140, 320M, 350M, 520M and 550M use the Cavium Nitrox Lite ASIC to accelerate AES, 3DES, 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 the CPU every seeing it. 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.

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 because of the large storage capabilities. 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.

2.2.2 ScreenOS

ScreenOS firmware 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 programming environment.

2.2.3 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. 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
- 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).

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 (Global policy lookup is not supported by the TOE in Authenticated Transparent Mode). The following sections describe differences between each of these three types of policies.

2.2.3.1 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.

2.2.3.2 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.

2.2.3.3 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.

2.2.3.3.1 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.

2.2.3.3.2 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.

2.2.4 Services

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.

2.3 TOE Configurations

The TOE supports a variety of 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.

2.3.1 Interface Modes

The TOE supports three types of interface modes. These interface modes include a Transparent Mode, NAT Mode, and Route Mode, each of which determines 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.

2.3.1.1 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.

The FDP_IFC.1(1), FDP_IFC.1(4), FDP_IFF.1(1), and FDP_IFF.1(4) security functional requirements specify the requirements for protecting information flows on a security appliance when it is configured in transparent mode.

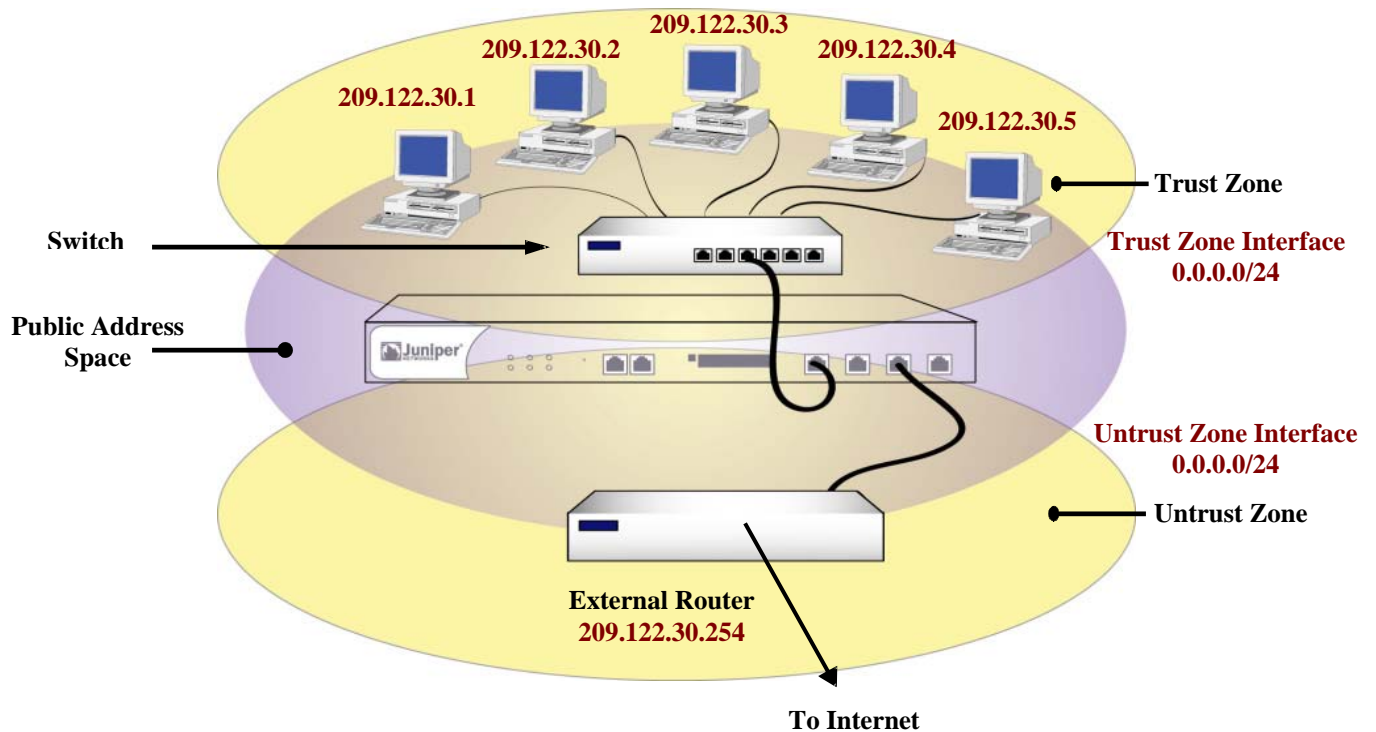


Figure 2.1: Transparent Mode

2.3.1.2 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.

The FDP_IFC.1(2), FDP_IFC.1(5), FDP_IFF.1(2) and FDP_IFF.1(5) security functional requirements specify the requirements for protecting information flows on a security appliance when it is configured in NAT mode.




Figure 2.2: NAT Mode

2.3.1.3 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.

The FDP_IFC.1(2), FDP_IFC.1(5), FDP_IFF.1(2) and FDP_IFF.1(5) security functional requirements specify the requirements for protecting information flows on a security appliance when it is configured in route mode.




Figure 2.3: Route Mode

2.3.2 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 Transparent Mode, and FDP_IFC.1(2) and FDP_IFF.1(2) for Route Mode and NAT Mode, for utilizing 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, 3DES, 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, 3DES, 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 all of the Security Function Policies (SFP) which include the TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, UNAUTHENTICATED TRANSPARENT MODE SFP, UNAUTHENTICATED ROUTE MODE SFP, and UNAUTHENTICATED TOE SERVICES SFP.

The TRANSPARENT MODE VPN SFP applies to traffic to or from a network interface configured in Transparent Mode that is using a VPN tunnel. The ROUTE MODE VPN SFP applies to traffic to or from a network interface configured in Route Mode or NAT Mode that is using a VPN tunnel. The UNAUTHENTICATED TRANSPARENT MODE SFP applies to traffic to or from a network interface configured in Transparent mode that is not using a VPN tunnel. The UNAUTHENTICATED ROUTE MODE SFP applies to traffic to or from a network interface configured in Route Mode or NAT Mode that is not using a VPN tunnel. the UNAUTHENTICATED TOE SERVICES SFP applies to traffic directed to the TOE.

2.3.2.1 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.

2.3.2.2 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 environment, provides the visual I/O for the administrative interface. The serial console is used in order to place the TOE into the evaluated configuration. Once the TOE has been placed into FIPS mode, the console may be used to monitor alarms, but is not to be used to enter commands. After the TOE is placed into the evaluated configuration, the administrative interface is provided over an SSH connection using encryption and certificate-based authentication.

The security appliance attaches to a physical network that has been separated into zones through port interfaces.

Security appliances come in several models. Each model differs in the performance capability, however all provide the same security functionality. 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 the security appliance, and from the security appliance to the receiving node in another network zone, if the security policy allows the information flow.

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.

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
- Resource utilization
- TOE access
- Trusted path/channels

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 or 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. More details about the audit mechanism can be found in Section 6.1.1, 'Security audit'.

2.4.2.2 Cryptographic support

The Juniper Networks Security Appliances are FIPS 140-2 validated as multi-chip standalone modules.

2.4.2.3 User data protection

The user data protection provided by the Security Appliance is provided through 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.

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

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.

2.4.2.3.2.1 Tunnel Interfaces

A tunnel interface acts as a doorway to a VPN tunnel. Traffic enters and exits a VPN tunnel via a tunnel interface.

When you bind a tunnel interface to a VPN tunnel, you can reference that tunnel interface in a route to a specific destination and then reference that destination in one or more policies. With this approach, you can finely control the flow of traffic through the tunnel. It also provides dynamic routing support for VPN traffic. When there is no tunnel interface bound to a VPN tunnel, you must specify the tunnel in the policy itself and choose tunnel as the action.

Outbound traffic enters the tunnel zone via the tunnel interface, is encapsulated, and exits via the security zone interface. Inbound traffic enters via the security zone interface, is decapsulated in the tunnel zone, and exits via the tunnel interface.

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 you connect to the security appliance via HTTP, SSH, or Telnet, you connect to the self zone. Remote management is supported in the evaluated configuration of the TOE via SSH.

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. FIPS 140-2 level 3 operator authentication requirements preclude the use of external authentication servers. Thus, to operate the TOE in a FIPS certified manner, only local administrator authentication is permitted in the evaluated configuration.

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.

More details about this mechanism can be found in Section 6.1.4, "Identification and authentication".

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. The web interface is not part of the evaluated configuration.

To execute the CLI, The administrator can establish a trusted SSH connection to the security appliance and utilize the CLI offered through the SSH connection. Regardless of the interface, 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 three distinct administrative roles: 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.

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 user communication with the security appliance itself. The TOE also preserves its configuration for a trusted recovery in the event that the configuration has been modified and not saved or if the security appliance has been ungracefully shutdown. The TOE additionally protects the session table by enforcing destination-based session limits and applying procedures to limit the lifetime of sessions when the session table reaches the defined watermark.

The TOE provides a recovery and self testing mechanism. The recovery mechanism allows administrators to return the TOE to a secure state, while the self test mechanism allows administrators to verify the integrity of the TOE and its cryptographic functions.

2.4.2.7 Resource utilization

The security appliance provides features to protect itself from Denial of Service attacks. These features limit TCP connections and offer administrators the ability to limit the number of resources a particular address or set of addresses can use over a specified time period.

2.4.2.8 TOE access

The security appliance provides the ability to restrict the establishment of an administrative session based on a schedule or based upon the originating source ip address (or subnet). The security appliance also provides inactivity timeouts and logon banners that can be configured by administrators.

2.4.2.9 Trusted path/channels

Remote administration of the security appliances can be accomplished using SSH to protect the communication of a remote administrator and the TOE. SSH provides for the protection of remote administration activity from both disclosure and modification. An IPSEC tunnel is used to provide encryption and integrity for trusted channels to external servers (e.g., an NTP 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 6.2.0 Concepts and Example, ScreenOS Reference Guide, Volume 1: Overview
- ScreenOS 6.2.0 Concepts and Example, ScreenOS Reference Guide, Volume 2: Fundamentals
- ScreenOS 6.2.0 Concepts and Example, ScreenOS Reference Guide, Volume 3: Administration
- ScreenOS 6.2.0 Concepts and Example, ScreenOS Reference Guide, Volume 4: Attack Detection
- ScreenOS 6.2.0 Concepts and Example, ScreenOS Reference Guide, Volume 5: VPNs
- ScreenOS 6.2.0 Concepts and Example, ScreenOS Reference Guide, Volume 8: Address Translation
- ScreenOS CLI Reference Guide: IPv6 Command Descriptions
- ScreenOS CLI Reference Guide: IPv4 Command Descriptions
- ScreenOS 6.2.0 Message Log Reference Guide
- Juniper Networks ScreenOS 6.2 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

3. Security Environment

All of the security environment statements have been drawn from a validated PP (Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile). Please consult those protection profiles for the description of the security environment. The policies, threats and assumptions from those PPs have been copied here for convenience. However, the VPN and TFFW PPs contain the definitive statement of security environment.

3.1 Organizational Policies

P.ACCESS_BANNER The TOE shall display an initial banner describing restrictions of use, legal agreements, or any other appropriate information to which users consent by accessing the system.

P.ACCOUNTABILITY The authorized users of the TOE shall be held accountable for their actions within the TOE.

P.ADMIN_ACCESS Administrators shall be able to administer the TOE both locally and remotely through protected communications channels.

P.CRYPTOGRAPHIC_FUNCTIONS The TOE shall provide cryptographic functions for its own use, including encryption/decryption and digital signature operations.

P.CRYPTOGRAPHY_VALIDATED Where the TOE requires FIPS-approved security functions, only NIST FIPS validated cryptography (methods and implementations) are acceptable for key management (i.e., generation, access, distribution, destruction, handling, and storage of keys) and cryptographic services (i.e., encryption, decryption, signature, hashing, key distribution, and random number generation services).

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).

P.VULNERABILITY_ANALYSIS_TEST The TOE must undergo appropriate independent vulnerability analysis and penetration testing to demonstrate that the TOE is resistant to an attacker possessing a medium attack potential.

3.2 Threats

T.ADDRESS_MASQUERADE A user on one interface may masquerade as a user on another interface to circumvent the TOE policy.

T.ADMIN_ERROR An administrator may incorrectly install or configure the TOE, or install a corrupted TOE resulting in ineffective security mechanisms.

T.ADMIN_ROGUE An administrator's intentions may become malicious resulting in user or TSF data being compromised.

T.AUDIT_COMPROMISE A malicious user or process may view audit records, cause audit records to be lost or modified, or prevent future audit records from being recorded, thus masking a user's action.

T.CRYPTO_COMPROMISE A malicious user or process may cause key, data or executable code associated with the cryptographic functionality to be inappropriately accessed (viewed, modified, or deleted), thus compromise the cryptographic mechanisms and the data protected by those mechanisms.

- T.FLAWED_DESIGN Unintentional or intentional errors in requirements specification or design of the TOE may occur, leading to flaws that may be exploited by a malicious user or program.
- T.FLAWED_IMPLEMENTATION Unintentional or intentional errors in implementation of the TOE design may occur, leading to flaws that may be exploited by a malicious user or program.
- T.MALICIOUS_TSF_COMPROMISE A malicious user or process may cause TSF data or executable code to be inappropriately accessed (viewed, modified, or deleted).
- T.MASQUERADE A user may masquerade as an authorized user or an authorized IT entity to gain access to data or TOE resources.
- T.POOR_TEST Lack of or insufficient tests to demonstrate that all TOE security functions operate correctly (including in a fielded TOE) may result in incorrect TOE behavior being undiscovered thereby causing potential security vulnerabilities.
- T.REPLAY A user may gain inappropriate access to the TOE by replaying authentication information, or may cause the TOE to be inappropriately configured by replaying TSF data or security attributes (captured as it was transmitted during the course of legitimate use).
- T.RESIDUAL_DATA A user or process may gain unauthorized access to data through reallocation of TOE resources from one user or process to another.
- T.RESOURCE_EXHAUSTION A malicious process or user may block others from TOE system resources (e.g., connection state tables) via a resource exhaustion denial of service attack.
- T.SPOOFING An entity may misrepresent itself as the TOE to obtain authentication data.
- T.UNATTENDED_SESSION A user may gain unauthorized access to an unattended session.
- T.UNAUTHORIZED_ACCESS A user may gain access to services (either on the TOE or by sending data through the TOE) for which they are not authorized according to the TOE security policy.
- T.UNAUTHORIZED_PEER An unauthorized IT entity may attempt to establish a security association with the TOE.
- T.UNIDENTIFIED_ACTIONS The administrator may fail to notice potential security violations, thus limiting the administrator's ability to identify and take action against a possible security breach.
- T.UNKNOWN_STATE When the TOE is initially started or restarted after a failure, design flaws, or improper configurations may cause the security state of the TOE to be unknown.

3.3 Assumptions

- A.NO_GENERAL_PURPOSE There are no general-purpose computing or storage repository capabilities (e.g., compilers, editors, or user applications) available on the TOE.
- A.NO_TOE_BYPASS Information cannot flow between external and internal networks located in different enclaves without passing through the TOE.
- A.PHYSICAL Physical security, commensurate with the value of the TOE and the data it contains, is assumed to be provided by the environment.

4. Security Objectives

All of the Security Objectives have been drawn from a validated PP (Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile). Please consult those protection profiles for the description of the security objectives. The security objectives from those PPs have been copied here for convenience. However, the VPN and TFFW PPs contain the definitive statement of security objectives.

4.1 Security Objectives for the TOE

- O.ADMIN_ROLE The TOE will provide administrator roles to isolate administrative actions, and to make the administrative functions available locally and remotely.
- O.AUDIT_GENERATION The TOE will provide the capability to detect and create records of security-relevant events associated with users.
- O.AUDIT_PROTECTION The TOE will provide the capability to protect audit information.
- O.AUDIT_REVIEW The TOE will provide the capability to selectively view audit information, and alert the administrator of identified potential security violations.
- O.CHANGE_MANAGEMENT The configuration of, and all changes to, the TOE and its development evidence will be analyzed, tracked, and controlled throughout the TOE's development.
- O.CORRECT_TSF_OPERATION The TOE will provide the capability to test the TSF to ensure the correct operation of the TSF in its operational environment.
- O.CRYPTOGRAPHIC_FUNCTIONS The TOE shall provide cryptographic functions for its own use, including encryption/decryption and digital signature operations.
- O.CRYPTOGRAPHY_VALIDATED The TOE shall use NIST FIPS 140-2 validated cryptomodules for cryptographic services implementing FIPS-approved security functions and random number generation services used by cryptographic functions.
- O.DISPLAY_BANNER The TOE will display an advisory warning regarding use of the TOE.
- O.DOCUMENT_KEY_LEAKAGE The bandwidth of channels that can be used to compromise key material shall be documented.
- O.INTEGRITY The TOE must be able to protect the integrity of data transmitted to a peer TOE 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.MAINT_MODE The TOE shall provide a mode from which recovery or initial startup procedures can be performed.
- O.MANAGE The TOE will provide all the functions and facilities necessary to support the administrators in their management of the security of the TOE, and restrict these functions and facilities from unauthorized use.
- O.MEDIATE The TOE must mediate the flow of information between sets of TOE network interfaces or between a network interface and the TOE itself in accordance with its security policy.

- O.PEER_AUTHENTICATION The TOE will authenticate each peer TOE that attempts to establish a security association with the TOE.
- O.REPLAY_DETECTION The TOE will provide a means to detect and reject the replay of TSF data and security attributes.
- O.RESIDUAL_INFORMATION The TOE will ensure that any information contained in a protected resource is not released when the resource is reallocated.
- O.RESOURCE_SHARING The TOE shall provide mechanisms that mitigate attempts to exhaust connection-oriented resources provided by the TOE (e.g., entries in a connection state table; Transmission Control Protocol (TCP) connections to the TOE).
- O.ROBUST_ADMIN_GUIDANCE The TOE will provide administrators with the necessary information for secure delivery and management.
- O.ROBUST_TOE_ACCESS The TOE will provide mechanisms that control a user's logical access to the TOE and to explicitly deny access to specific users when appropriate.
- O.SELF_PROTECTION The TSF will maintain a domain for its own execution that protects itself and its resources from external interference, tampering, or unauthorized disclosure.
- O.SOUND_DESIGN The design of the TOE will be the result of sound design principles and techniques; the design of the TOE, as well as the design principles and techniques, are adequately and accurately documented.
- O.SOUND_IMPLEMENTATION The implementation of the TOE will be an accurate instantiation of its design, and is adequately and accurately documented.
- O.THOROUGH_FUNCTIONAL_TESTING The TOE will undergo appropriate security functional testing that demonstrates the TSF satisfies the security functional requirements.
- O.TIME_STAMPS The TOE shall provide reliable time stamps and the capability for the administrator to set the time used for these time stamps.
- O.TRUSTED_PATH The TOE will provide a means to ensure users are not communicating with some other entity pretending to be the TOE, and that the TOE is communicating with an authorized IT entity and not some other entity pretending to be an authorized IT entity.
- O.VULNERABILITY_ANALYSIS_TEST The TOE will undergo appropriate independent vulnerability analysis and penetration testing to demonstrate the design and implementation of the TOE does not allow attackers with medium attack potential to violate the TOE's security policies.

4.2 Security Objectives for the Environment

- OE.CRYPTANALYTIC Cryptographic methods used in the IT environment shall be interoperable with the TOE, should be FIPS 140-2 validated and should be resistant to cryptanalytic attacks (i.e., will be of adequate strength to protect unclassified Mission Support, Administrative, or Mission Critical data).
- OE.NO_GENERAL_PURPOSE The Administrator ensures there are no general-purpose computing or storage repository capabilities (e.g., compilers, editors, or user applications) available on the TOE.
- OE.NO_TOE_BYPASS Information cannot flow between external and internal networks located in different enclaves without passing through the TOE.

OE.PHYSICAL Physical security, commensurate with the value of the TOE and the data it contains, is assumed to be provided by the IT environment.

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 Medium Robustness VPN Protection Profile and Medium Robustness TFW Protection Profile. 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.

This security target utilizes extended requirements only as reproductions of requirements found in protection profiles to which this security target is claiming compliance. Therefore, all requirements for information related to the extended requirements is satisfied by this security target's compliance with validated protection profiles.

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_ARP.1: Security alarms
	FAU_ARP_ACK_(EXT).1: Security alarm acknowledgement
	FAU_GEN.1-NIAP-0410: Audit data generation
	FAU_GEN.2-NIAP-0410: User identity association
	FAU_SAA.1-NIAP-0407: Potential violation analysis
	FAU_SAR.1: Audit review
	FAU_SAR.2: Restricted audit review
	FAU_SAR.3: Selectable audit review
	FAU_SEL.1-NIAP-0407: Selective Audit
	FAU_STG.1-NIAP-0429: Protected audit trail storage
	FAU_STG.3: Action in case of possible audit data loss
	FAU_STG.NIAP-0414-1-NIAP-0429: Site-Configurable Prevention of Audit Loss
	FCS: Cryptographic support
FCS_CKM.1(1): Cryptographic key generation (for symmetric keys)	
FCS_CKM.1(2): Cryptographic key generation (for asymmetric keys)	
FCS_CKM.2: Cryptographic key distribution	
FCS_CKM_(EXT).2: Cryptographic Key Handling and Storage	
FCS_CKM.4: Cryptographic key destruction	
FCS_COP.1(1): Cryptographic Operation (for data encryption/decryption)	
FCS_COP.1(2): Cryptographic operation (for	

	cryptographic signature)
	FCS_COP.1(3): Cryptographic operation (for cryptographic hashing)
	FCS_COP.1(4): Cryptographic operation (for cryptographic key agreement)
	FCS_COP_(EXT).1: Random Number Generation
	FCS_IKE_(EXT).1: Internet Key Exchange
FDP: User data protection	FDP_IFC.1(1): Subset information flow control (VPN Policy-Transparent Mode)
	FDP_IFC.1(2): Subset information flow control (VPN Policy-Route Mode)
	FDP_IFC.1(3): Subset information flow control (Unauthenticated TOE Services Policy)
	FDP_IFC.1(4): Subset information flow control (Transparent Mode Firewall Policy)
	FDP_IFC.1(5): Subset information flow control (Route Mode Firewall Policy)
	FDP_IFF.1(1): Simple security attributes (VPN Policy-Transparent Mode)
	FDP_IFF.1(2): Simple security attributes (VPN Policy-Route Mode)
	FDP_IFF.1(3): Simple security attributes (Unauthenticated TOE Services Policy)
	FDP_IFF.1(4): Simple security attributes (Transparent Mode Firewall Policy)
	FDP_IFF.1(5): Simple security attributes (Route Mode Firewall Policy)
	FDP_RIP.2: Full residual information protection
FIA: Identification and authentication	FIA_AFL.1: Authentication failure handling
	FIA_ATD.1: User attribute definition
	FIA_UAU.1: Timing of authentication
	FIA_UAU.2: User authentication before any action
	FIA_UAU_(EXT).5: Multiple authentication mechanisms
	FIA_UID.2: User identification before any action
	FIA_USB.1: User-subject binding
FMT: Security management	FMT_MOF.1(1): Management of security functions behavior (TSF non-cryptographic self-test)
	FMT_MOF.1(2): Management of security functions behavior (cryptographic self-test)
	FMT_MOF.1(3): Management of security functions behavior (audit and alarms)
	FMT_MOF.1(4): Management of security functions behavior (audit and alarms)
	FMT_MOF.1(5): Management of security functions behavior (audit and alarms)
	FMT_MOF.1(6): Management of security functions behavior (available TOE-services for unauthenticated users)

	FMT_MOF.1(7): Management of security functions behavior (quota mechanism)
	FMT_MSA.1: Management of security attributes
	FMT_MSA.3(1): Static attribute initialization (attributes/ruleset)
	FMT_MSA.3(2): Static attribute initialization (services)
	FMT_MTD.1(1): Management of TSF data (non-cryptographic, non-time TSF data)
	FMT_MTD.1(2): Management of TSF data (cryptographic TSF data)
	FMT_MTD.1(3): Management of TSF data (time TSF data)
	FMT_MTD.1(4): Management of TSF data (Policy Rulesets)
	FMT_MTD.2(1): Management of limits on TSF data (transport-layer quotas)
	FMT_MTD.2(2): Management of limits on TSF data (controlled connection-oriented quotas)
	FMT_REV.1: Revocation
	FMT_SMF.1: Specification of Management Functions
	FMT_SMR.2: Restrictions on security roles
FPT: Protection of the TSF	FPT_RCV.1: Manual recovery
	FPT_RPL.1: Replay detection
	FPT_STM.1: Reliable time stamps
	FPT_TST_(EXT).1: TSF testing
	FPT_TST.1(1): TSF Testing (for cryptography)
	FPT_TST.1(2): TSF Testing (for key generation components)
FRU: Resource utilization	FRU_RSA.1(1): Maximum quotas (transport-layer quotas)
	FRU_RSA.1(2): Maximum quotas (controlled connection-oriented quotas)
FTA: TOE access	FTA_SSL.1: TSF-initiated session locking
	FTA_SSL.2: User-initiated locking
	FTA_SSL.3: TSF-initiated termination
	FTA_TAB.1: Default TOE access banners
	FTA_TSE.1: TOE session establishment
FTP: Trusted path/channels	FTP_ITC.1(1): Inter-TSF trusted channel (Prevention of Disclosure)
	FTP_ITC.1(2): Inter-TSF trusted channel (Prevention of Modification)
	FTP_TRP.1(1): Trusted Path (Prevention of Disclosure)
	FTP_TRP.1(2): Trusted Path (Detection of Modification)

Table 1 TOE Security Functional Requirements**5.1.1 Security audit (FAU)****5.1.1.1 Security alarms (FAU_ARP.1)**

FAU_ARP.1.1 The TSF shall [immediately display an alarm message, identifying the potential security violation and make accessible the audit record contents associated with the auditable event(s) that generated the alarm, at the:

- a) local console,
- b) remote administrator sessions that exist, and;
- c) remote administrator sessions that are initiated before the alarm has been acknowledged, and;
- d) at the option of the Security Administrator, generate an audible alarm, and;
- e) [*no other methods*]

upon detection of a potential security violation.

5.1.1.2 Security alarm acknowledgement (FAU_ARP_ACK_(EXT).1)

FAU_ARP_ACK_(EXT).1.1 The TSF shall display the alarm message identifying the potential security violation and make accessible the audit record contents associated with the auditable event(s) until it has been acknowledged. An audible alarm will sound until acknowledged by an administrator.

FAU_ARP_ACK_(EXT).1.2 The TSF shall display an acknowledgement message identifying a reference to the potential security violation, a notice that it has been acknowledged, the time of the acknowledgement and the user identifier that acknowledged the alarm, at the:

- a) local console, and
- b) remote administrator sessions that received the alarm.

5.1.1.3 Audit data generation (FAU_GEN.1-NIAP-0410)

FAU_GEN.1-NIAP-0410.1 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 auditable events [listed in ~~Table 3~~ **Section 9 of this security target**];
- c) [*no additional events*].

FAU_GEN.1-NIAP-0410.2 The TSF shall record within each audit record at least the following information:

- a) Date and time of the event, type of event, and the 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 three of ~~Table 7 below~~ **Section 9 of this security target**].

5.1.1.4 User identity association (FAU_GEN.2-NIAP-0410)

FAU_GEN.2-NIAP-0410.1 The TSF shall be able to associate each auditable event with the identity of the user that caused the event.

5.1.1.5 Potential violation analysis (FAU_SAA.1-NIAP-0407)

FAU_SAA.1-NIAP-0407.1 The TSF shall be able to apply a set of rules in monitoring events and based upon these rules indicate a potential violation of the TSP.

FAU_SAA.1-NIAP-0407.2 The TSF shall enforce the following rules for monitoring audited events:

- a) [Security Administrator specified number of authentication failures;
- b) Security Administrator specified number of Information Flow policy violations by an individual presumed source network identifier (e.g., IP address) within an administrator specified time period;
- c) Security Administrator specified number of Information Flow policy violations to an individual destination network identifier within an administrator specified time period;
- d) Security Administrator specified number of Information Flow policy violations to an individual destination subject service identifier (e.g., TCP port) within an administrator specified time period;
- e) Security Administrator specified Information Flow policy rule, or group of rule violations within an administrator specified time period;
- f) Any detected replay of TSF data or security attributes;
- g) Any failure of the cryptomodule self-tests (FPT_TST.1(1));
- h) Any failure of the other key generation self-tests (FPT_TST.1(2));
- i) Any failure of the other TSF self-tests (FPT_TST_(EXT).1);
- j) Security Administrator specified number of encryption failures;
- k) Security Administrator specified number of decryption failures;
- l) Security Administrator specified number of Phase 1 authentication failures when negotiating the Internet Key Exchange protocol;
- m) Security Administrator specified number of failures occur during Phase 2 negotiation; and
- n) [*no additional rules*]

known to indicate a potential security violation;

5.1.1.6 Audit review (FAU_SAR.1)

FAU_SAR.1.1 The TSF shall provide [the administrators] with the capability to read [all audit data] from the audit records.

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

5.1.1.7 Restricted audit review (FAU_SAR.2)

FAU_SAR.2.1 The TSF shall prohibit all users read access to the audit records in the audit trail, except the Administrators.

5.1.1.8 Selectable audit review (FAU_SAR.3)

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

- a) [user identity;
- b) source subject identity;
- c) destination subject identity;
- d) ranges of one or more: dates, times, user identities, subject service identifiers, or transport layer protocol;
- e) TOE network interfaces;
- f) [*rule identity*;]

5.1.1.9 Selective Audit (FAU_SEL.1-NIAP-0407)

FAU_SEL.1-NIAP-0407.1 The TSF shall allow only the Security Administrator to include or exclude auditable events from the set of audited events based on the following attributes:

- a) user identity;
- b) event type;
- c) [network identifier;
- d) subject service identifier;
- e) success of auditable security events;
- f) failure of auditable security events;
- g) [*rule identity*];].

5.1.1.10 Protected audit trail storage (FAU_STG.1-NIAP-0429)

FAU_STG.1-NIAP-0429.1 The TSF shall restrict the deletion of stored audit records in the audit trail to the Audit Administrator.

FAU_STG.1-NIAP-0429.2 The TSF shall be able to prevent (unauthorized) modifications to the audit records in the audit trail.

5.1.1.11 Action in case of possible audit data loss (FAU_STG.3)

FAU_STG.3.1 The TSF shall [immediately alert the administrators by displaying a message at the local console, and at the remote administrative console when an administrative session exists for each of the defined administrative roles, at the option of the Security Administrator generate an audible alarm, [*no other methods*]] if the audit trail exceeds [a Security Administrator settable percentage of storage capacity].

5.1.1.12 Site-Configurable Prevention of Audit Loss (FAU_STG.NIAP-0414-1-NIAP-0429)

FAU_STG.NIAP-0414-1-NIAP-0429.1 The TSF shall provide the Security Administrator the capability to select one or more of the following actions prevent auditable events, except those taken by the Security Administrator and Audit Administrator, overwrite the oldest stored audit records and [*no other actions*] to be taken if the audit trail is full.

FAU_STG.NIAP-0414-1-NIAP-0429.2 The TSF shall enforce the Security Administrator's selection(s) if the audit trail is full.

5.1.2 Cryptographic support (FCS)

5.1.2.1 Baseline Cryptographic Module (FCS_BCM_(EXT).1)

FCS_BCM_(EXT).1.1 All FIPS-approved cryptographic functions implemented by the TOE shall be implemented in a cryptomodule that is FIPS 140-2 validated, and perform the specified cryptographic functions in a FIPS-approved mode of operation. The FIPS 140-2 validation shall include an algorithm validation certificate for all FIPS-approved cryptographic functions implemented by the TOE.

FCS_BCM_(EXT).1.2 All cryptographic modules implemented in the TSF [*as a combination of hardware and software shall have a minimum overall rating of FIPS PUB 140-2, Level 1 and also meet FIPS PUB 140-2, Level 3 for the following:*

- *Cryptographic Module Ports and Interfaces;*
- *Roles, Services and Authentication;*
- *Cryptographic Key Management; and*
- *Design Assurance].*

5.1.2.2 Cryptographic key generation (for symmetric keys) (FCS_CKM.1(1))

FCS_CKM.1.1(1) The TSF shall generate symmetric cryptographic keys using a FIPS-Approved Random Number Generator as specified in FCS_COP_(EXT).1, and provide integrity protection to generated symmetric keys in accordance with NIST SP 800-57 “Recommendation for Key Management” Section 6.1.

5.1.2.3 Cryptographic key generation (for asymmetric keys) (FCS_CKM.1(2))

FCS_CKM.1.1(2) The TSF shall generate asymmetric cryptographic keys in accordance with the mathematical specifications of the FIPS-approved or NIST-recommended standard [ANSI X9.62-1998], using a domain parameter generator and [

- (1) a FIPS-Approved Random Number Generator as specified in FCS_COP_(EXT).1, and/or
- (2) a prime number generator as specified in ANSI X9.80 “Prime Number Generation, Primality Testing, and Primality Certificates” using random integers with deterministic tests, or constructive generation methods]

in a cryptographic key generation scheme that meets the following:

- The TSF shall provide integrity protection and assurance of domain parameter and public key validity to generated asymmetric keys in accordance with NIST SP 800-57 “Recommendation for Key Management” Section 6.1.
- Generated key strength shall be equivalent to, or greater than, a symmetric key strength of 128 bits using conservative estimates.

5.1.2.4 Cryptographic key distribution (FCS_CKM.2)

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a specified cryptographic key distribution method [**Manual (Physical) Method and Automated (electronic) Method**] that meets the following: [

- NIST Special Publication 800-57, “Recommendation for Key Management” Section 8.1.5.
- NIST Special Publication 800-56A, “Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography”].

5.1.2.5 Cryptographic Key Handling and Storage (FCS_CKM_(EXT).2)

FCS_CKM_(EXT).2.1 The TSF shall perform a key error detection check on each transfer of key (internal, intermediate transfers).

FCS_CKM_(EXT).2.2 The TSF shall store persistent secret and private keys when not in use in encrypted form or using split knowledge procedures.

FCS_CKM_(EXT).2.3 The TSF shall destroy non-persistent cryptographic keys after a cryptographic administrator-defined period of time of inactivity.

FCS_CKM_(EXT).2.4 The TSF shall prevent archiving of expired (private) signature keys.

5.1.2.6 Cryptographic key destruction (FCS_CKM.4)

FCS_CKM.4.1 The TSF shall destroy cryptographic keys in accordance with a cryptographic key zeroization method that meets the following:

- a) the key zeroization requirements in FIPS PUB 140-2, “Security Requirements for Cryptographic Modules”;
- b) Zeroization of all plaintext cryptographic keys and all other critical cryptographic security parameters shall be immediate and complete; and

- c) The TSF shall zeroize each intermediate storage area for plaintext key/critical cryptographic security parameter (i.e., any storage, such as memory buffers, that is included in the path of such data) upon the transfer of the key/critical cryptographic security parameter to another location..
- d) For non-volatile memories other than EEPROM and Flash, the zeroization shall be executed by overwriting three or more times using a different alternating data pattern each time.
- e) For volatile memory and non-volatile EEPROM and Flash memories, the zeroization shall be executed by a single direct overwrite consisting of a pseudo random pattern, followed by a read-verify..

5.1.2.7 Cryptographic Operation (for data encryption/decryption) (FCS_COP.1(1))

FCS_COP.1.1(1) A cryptomodule shall perform encryption and decryption using the FIPS-Approved Security Function AES algorithm operating in [*CBC*] mode(s) supporting key sizes of [*128 bits, 192 bits, and 256 bits*].

5.1.2.8 Cryptographic operation (for cryptographic signature) (FCS_COP.1(2))

FCS_COP.1.1(2) The TSF shall perform cryptographic signature services using the FIPS-approved security function [

- a) Digital Signature Algorithm (DSA) with a key size (modulus) of [**2048 bits or greater**],
- b) RSA Digital Signature Algorithm (rDSA) with a key size (modulus) of [**2048 bits or greater**], or
- c) Elliptic Curve Digital Signature Algorithm (ECDSA) with a key size of [**256 bits**], using only the NIST curve(s) [**P-256**]]

that meets NIST Special Publication 800-57, “Recommendation for Key Management.”

5.1.2.9 Cryptographic operation (for cryptographic hashing) (FCS_COP.1(3))

FCS_COP.1.1(3) The TSF shall perform cryptographic hashing services using the FIPS-approved security function Secure Hash Algorithm and any message digest specified in FIPS 180-2 [SHA256].

5.1.2.10 Cryptographic operation (for cryptographic key agreement) (FCS_COP.1(4))

FCS_COP.1.1(4) The TSF shall perform cryptographic key agreement services using the FIPS-approved security function as specified in NIST Special Publication 800-56A, “Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography” [

- (1) [Finite Field-based key agreement algorithm] and cryptographic key sizes(modulus) of [2048 bits or greater], or
- (2) [Elliptic Curve-based key agreement algorithm] and cryptographic key size of [256 bits or greater] using only the NIST curve(s) [P-256]]

that meets NIST Special Publication 800-57, “Recommendation for key Management”.

5.1.2.11 Random Number Generation (FCS_COP_(EXT).1)

FCS_COP_(EXT).1.1 The TSF shall perform all random number generation (RNG) services in accordance with a FIPS-approved RNG [**ANSI X9.31 algorithm**] seeded by [

- (1) one or more independent hardware-based entropy sources, and/or
- (2) one or more independent software-based entropy sources, and/or
- (3) a combination of hardware-based and software-based entropy sources.]

FCS_COP_(EXT).1.2 The TSF shall defend against tampering of the random number generation (RNG)/ pseudorandom number generation (PRNG) sources.

5.1.2.12 Internet Key Exchange (FCS_IKE_(EXT).1)

FCS_IKE_(EXT).1.1 The TSF shall provide cryptographic key establishment techniques in accordance with RFC 2409 as follows(s):

- Phase 1, the establishment of a secure authenticated channel between the TOE and another remote VPN endpoint, shall be performed using one of the following, as configured by the security administrator:
 - Main Mode
 - Aggressive Mode
- Phase 2, negotiation of security services for IPsec, shall be done using Quick Mode, using SHA-1 as the pseudo-random function. Quick Mode shall generate key material that provides perfect forward secrecy. The use of SHA-256 and SHA-384 as the PRF in IKE1 KDF is also allowed.

FCS_IKE_(EXT).1.2 The TSF shall require the x of g^{xy} be randomly generated using a FIPS-approved random number generator when computation is being performed. The minimum size of x shall be twice the number of bits of the strength level associated with the negotiated DH group per table 2 of NIST SP 800-57. The nonce sizes are to be between 8 and 256 bytes. Nonces shall be generated in a manner such that the probability that a specific nonce value will be repeated during the life a specific IPsec SA is less than 1 in $2^{(\text{bit strength of the negotiated DH group})}$

FCS_IKE_(EXT).1.3 When performing authentication using pre-shared keys, the key shall be generated using the FIPS approved random number generator specified in FCS_COP_(EXT).1.1.

FCS_IKE_(EXT).1.4 The TSF shall compute the value of SKEYID (as defined in RFC 2409), using SHA-1 as the pseudo-random function. The use of SHA-256 and SHA-384 as the PRF in IKEv1 KDF is also allowed. The TSF shall be capable of authentication using the methods for

- Signatures: $\text{SKEYID} = \text{sha}(\text{Ni}_b \mid \text{Nr}_b, g^{xy})$
- Pre-shared keys: $\text{SKEYID} = \text{sha}(\text{pre-shared-key}, \text{Ni}_b \mid \text{Nr}_b)$
- **[Authentication using Public key encryption, computing SKEYID as follows:**
 $\text{SKEYID} = \text{sha}(\text{sha}(\text{Ni}_b \mid \text{Nr}_b), \text{CKY-I} \mid \text{CKY-R})$]

FCS_IKE_(EXT).1.5 The TSF shall compute authenticated keying material as follows:

- $\text{SKEYID}_d = \text{sha}(\text{SKEYID}, g^{xy} \mid \text{CKY-I} \mid \text{CKY-R} \mid 0)$
- $\text{SKEYID}_a = \text{sha}(\text{SKEYID}, \text{SKEYID}_d \mid g^{xy} \mid \text{CKY-I} \mid \text{CKY-R} \mid 1)$
- $\text{SKEYID}_e = \text{sha}(\text{SKEYID}, \text{SKEYID}_a \mid g^{xy} \mid \text{CKY-I} \mid \text{CKY-R} \mid 2)$
- [[none]]

FCS_IKE_(EXT).1.6 To authenticate the Phase 1 exchange, the TSF shall generate HASH_I if it is the initiator, or HASH_R if it is the responder as follows:

- $\text{HASH}_I = \text{sha}(\text{SKEYID}, g^{xi} \mid g^{xr} \mid \text{CKY-I} \mid \text{CKY-R} \mid \text{SAi}_b \mid \text{IDii}_b)$
- $\text{HASH}_R = \text{sha}(\text{SKEYID}, g^{xr} \mid g^{xi} \mid \text{CKY-R} \mid \text{CKY-I} \mid \text{SAi}_b \mid \text{IDir}_b)$

FCS_IKE_(EXT).1.7 The TSF shall be capable of authenticating IKE Phase 1 using the following methods as defined in RFC 2409, as configured by the security administrator:

- a) Authentication with digital signatures: The TSF shall use [RSA, DSA, **[and ECDSA]**]
- b) when an RSA signature is applied to HASH I or HASH R it must be first PKCS#1 encoded. The TSF shall check the HASH_I and HASH_R values sent against a computed value to detect any changes made to the proposed transform negotiated in phase one. If changes are detected the session shall be terminated and an alarm shall be generated.
- c) **[/X.509 certificates Version 3 and no other versions:]** X.509 V3 implementations, if implemented, shall be capable of checking for validity of the certificate path, and at option of SA, check for certificate revocation.
- d) Authentication with a pre-shared key: The TSF shall allow authentication using a pre-shared key.

FCS_IKE_(EXT).1.8 The TSF shall compute the hash values for Quick Mode in the following way
 $\text{HASH}(1) = \text{sha}(\text{SKEYID}_a, \text{M-ID} \mid [\text{any ISAKMP payload after HASH}(1) \text{ header contained in the message}])$

$\text{HASH}(2) = \text{sha}(\text{SKEYID}_a, \text{M-ID} \mid \text{Ni}_b \mid [\text{any ISAKMP payload after HASH}(2) \text{ header contained in the message}])$

$\text{HASH}(3) = \text{sha}(\text{SKEYID}_a, 0 \mid \text{M-ID} \mid \text{Ni}_b \mid \text{Nr}_b)$

FCS_IKE_(EXT).1.9 The TSF shall compute new keying material during Quick Mode as follows:

[selection: when using perfect forward secrecy
 KEYMAT = sha(SKEYID_d, g(qm)^xy | protocol | SPI | Ni_b | Nr_b),
 When perfect forward secrecy is not used
 KEYMAT = sha(SKEYID_d | protocol | SPI | Ni_b | Nr_b)]

FCS_IKE_(EXT).1.10 The TSF shall at a minimum, support the following ID types: [ID_IPV4_ADDR, ID_FQDN, ID_USER_FQDN, ID_IPV6_ADDR, [*ID_IPV4_ADDR_SUBNET*, *ID_IPV6_ADDR_SUBNET*]].

5.1.3 User data protection (FDP)

5.1.3.1 Subset information flow control (VPN Policy-Transparent Mode) (FDP_IFC.1(1))

FDP_IFC.1.1(1) The TSF shall enforce the [**TRANSPARENT MODE VPN SFP**] on [

- a) source subject: TOE interface on which information is received;
- b) destination subject: TOE interface to which information is destined.
- c) information: network packets; and
- d) 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**].

5.1.3.2 Subset information flow control (VPN Policy-Route Mode) (FDP_IFC.1(2))

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

- a) source subject: TOE interface on which information is received;
- b) destination subject: TOE interface to which information is destined.
- c) information: network packets; and
- e) 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 Subset information flow control (Unauthenticated TOE Services Policy) (FDP_IFC.1(3))

FDP_IFC.1.1(3) The TSF shall enforce the [**UNAUTHENTICATED TOE SERVICES SFP**] on [

- a) source subject: TOE interface on which information is received;
- b) destination subject: the TOE;
- c) information: network packets; and
- d) operations: accept or reject network packet].

5.1.3.4 Subset information flow control (Transparent Mode Firewall Policy) (FDP_IFC.1(4))

FDP_IFC.1.1(4) The TSF shall enforce the [**TRANSPARENT MODE FIREWALL SFP**] on [

- a) source subject: TOE interface on which information is received;
- b) destination subject: TOE interface on which information is destined;
- c) information: network packets; and
- d) operations: pass information].

5.1.3.5 Subset information flow control (Route Mode Firewall Policy) (FDP_IFC.1(5))

- FDP_IFC.1.1(5)** The TSF shall enforce the [**ROUTE MODE FIREWALL SFP**] on [
- a) source subject: TOE interface on which information is received;
 - b) destination subject: OE interface on which information is destined;
 - c) information: network packets; and
 - d) operations: pass information].

5.1.3.6 Simple security attributes (VPN Policy-Transparent Mode) (FDP_IFF.1(1))

- FDP_IFF.1.1(1)** The TSF shall enforce the [**TRANSPARENT MODE VPN SFP**] based on the following types of subject and information security attributes:

- a) [Source subject security attributes:
 - set of source subject identifiers; and
 - [*none*].
- b) Destination subject security attributes:
 - Set of destination subject identifiers; and
 - [*none*].
- c) Information security attributes:
 - presumed identity of source subject;
 - identity of destination subject;]

- FDP_IFF.1.2(1)** 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
 - [**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.**
 - **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(1) 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 [the Security Administrator shall have the capability to view all information flows allowed by the information flow policy ruleset before the ruleset is applied].

- 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 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.).

5.1.3.7 Simple security attributes (VPN Policy-Route Mode) (FDP_IFF.1(2))

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

- a) [Source subject security attributes:
 - set of source subject identifiers; and
 - **[none]**.
- b) Destination subject security attributes:
 - Set of destination subject identifiers; and
 - **[none]**.
- 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
 - **[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.**
 - **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(2) The TSF shall enforce the **[none]**.

FDP_IFF.1.4(2) The TSF shall provide the following [

- a) the Security Administrator shall have the capability to view all information flows allowed by the information flow policy ruleset before the ruleset is applied;
- b) **the TSF shall provide the capability to perform policy-based address translation on the presumed IPv4 address of the source subject, in the information, if a policy with one of the following policy-based translation options is invoked by the information:**
 - **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**
- c) **the TSF shall provide the capability to perform policy-based address translation on the presumed IPv4 address of the destination subject, in the information, if a policy with one of the following policy-based translation options is invoked by the information:**
 - **NAT-Dst to a Single IPv4 Address with Port Mapping**

- **NAT-Dst to a Single IPv4 Address without Port Mapping**
- **NAT-Dst from an IPv4 Address Range to a Single IP Address**
- **NAT-Dst between IPv4 Address Ranges].**

FDP_IFF.1.5(2) 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.).

5.1.3.8 Simple security attributes (Unauthenticated TOE Services Policy) (FDP_IFF.1(3))

FDP_IFF.1.1(3) The TSF shall enforce the [UNAUTHENTICATED TOE SERVICES SFP] based on the following types of subject and information security attributes:

- a) [Source subject security attributes:
 - set of source subject identifiers; and
 - [none].
- b) Destination subject security attributes:
 - TOE's network identifier; and
 - [none].
- c) Information security attributes:
 - presumed identity of source subject;
 - identity of destination subject;
 - transport layer protocol;
 - source subject service identifier;
 - destination subject service identifier (e.g., TCP or UDP destination port number); and
 - [ICMP echo and ARP communications].

FDP_IFF.1.2(3) The TSF shall permit an information flow between a source subject and the TOE 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 the TOE;
- the information security attributes match the attributes in an information flow control policy according to the following algorithm
 - [Packets are accepted if and only if they are for one of the TOE services that are explicitly configured as permitted for the interface on which the packets arrive; and
 - The packet's source address must match an address in the set of 'Manager-IP' addresses for services except ICMP ping;
 - otherwise the packet is rejected]].

FDP_IFF.1.3(3) The TSF shall enforce the [following rules:

- The TOE shall allow source subjects to access TOE services [ICMP echo] without authenticating those source subjects; and

- The TOE shall allow the list of services specified immediately above to be enabled (become available to unauthenticated users) or disabled (become unavailable to unauthenticated users)].

FDP_IFF.1.4(3) The TSF shall provide the following [the Security Administrator shall have the capability to view all information flows allowed by this information flow control policy before the policy is applied].

FDP_IFF.1.5(3) The TSF shall explicitly authorize an information flow based on the following rules: [

- ***Unauthenticated ARP communications received by the TOE are allowed from all subjects***].

FDP_IFF.1.6(3) 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; and
- 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 TOE].

5.1.3.9 Simple security attributes (Transparent Mode Firewall Policy) (FDP_IFF.1(4))

FDP_IFF.1.1(4) The TSF shall enforce the [TRANSPARENT MODE FIREWALL SFP] based on the following types of subject and information security attributes:

- a) [Source subject security attributes:
 - set of source subject identifiers; and
 - *[none]*.
- b) Destination subject security attributes:
 - TOE's network identifier; and
 - *[none]*.
- c) Information security attributes:
 - presumed identity of source subject;
 - identity of destination subject;
 - transport layer protocol;
 - source subject service identifier;
 - destination subject service identifier (e.g., TCP or UDP destination port number); and
 - *[none]*
 - Stateful packet attributes: [
 - Connection-oriented protocols:
 - sequence number;
 - acknowledgement number;
 - Flags:
 - SYN,
 - ACK,
 - RST,
 - FIN; and
 - *none*.
 - Connectionless protocols:
 - source and destination network identifiers;
 - source and destination service identifiers;
 - *none*]].

FDP_IFF.1.2(4) 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
 - **[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.**
 - **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].

FDP_IFF.1.3(4) The TSF shall enforce the [following:

- fragmentation rule:
 - prior to applying the information policy ruleset, the TOE completely reassembles fragmented packets;
- stateful packet inspection rules:
 - 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, as defined in FDP_IFF.1.2(4), is applied to the packet;
 - otherwise, the TSF associates a packet with an allowed established session using the stateful packet attributes].

FDP_IFF.1.4(4) The TSF shall provide the following [the Security Administrator shall have the capability to view all information flows allowed by this information flow control policy before the policy is applied].

FDP_IFF.1.5(4) The TSF shall explicitly authorize an information flow based on the following rules: [none].

FDP_IFF.1.6(4) 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; and
- 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 TOE].

5.1.3.10 Simple security attributes (Route Mode Firewall Policy) (FDP_IFF.1(5))

FDP_IFF.1.1(5) The TSF shall enforce the [**ROUTE MODE FIREWALL SFP**] based on the following types of subject and information security attributes:

- a) [Source subject security attributes:
 - set of source subject identifiers; and
 - *[none]*.
- b) Destination subject security attributes:
 - Set of destination subject identifiers; and
 - *[none]*.

- c) Information security attributes:
- presumed identity of source subject;
 - identity of destination subject;
 - transport layer protocol;
 - source subject service identifier;
 - destination subject service identifier (e.g., TCP or UDP destination port number); and
 - *[none]*
 - Stateful packet attributes: [
 - Connection-oriented protocols:
 - sequence number;
 - acknowledgement number;
 - Flags:
 - SYN,
 - ACK,
 - RST,
 - FIN; and
 - *none*.
 - Connectionless protocols:
 - source and destination network identifiers;
 - source and destination service identifiers;
 - *none*].

FDP_IFF.1.2(5) 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
 - **[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.**
 - **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].

FDP_IFF.1.3(5) The TSF shall enforce the [following:

- fragmentation rule:
 - prior to applying the information policy ruleset, the TOE completely reassembles fragmented packets;
- stateful packet inspection rules:
 - 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, as defined in FDP_IFF.1.2(5), is applied to the packet;
 - otherwise, the TSF associates a packet with an allowed established session using the stateful packet attributes].

- FDP_IFF.1.4(5)** The TSF shall provide the following [
- a) the Security Administrator shall have the capability to view all information flows allowed by the information flow policy ruleset before the ruleset is applied;
 - b) **the TSF shall provide the capability to perform policy-based address translation on the presumed IPv4 address of the source subject, in the information, if a policy with one of the following policy-based translation options is invoked by the information:**
 - 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
 - c) **the TSF shall provide the capability to perform policy-based address translation on the presumed IPv4 address of the destination subject, in the information, if a policy with one of the following policy-based translation options is invoked by the information:**
 - NAT-Dst to a Single IPv4 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].

FDP_IFF.1.5(5) The TSF shall explicitly authorize an information flow based on the following rules: [none].

- FDP_IFF.1.6(5)** 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; and
 - 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 TOE].

5.1.3.11 Full residual information protection (FDP_RIP.2)

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 Authentication failure handling (FIA_AFL.1)

FIA_AFL.1.1 The TSF shall detect when [a Security Administrator-configurable positive integer] of unsuccessful authentication attempts occur related to [administrators attempting to authenticate remotely and authorized IT entities].

FIA_AFL.1.2 When the defined number of unsuccessful authentication attempts has been met, the TSF shall [at the option of the Security Administrator prevent the remote administrators, or an authorized IT entity from

1. performing activities that require authentication until an action is taken by the Security Administrator, or
2. until a Security Administrator defined time period has elapsed].

5.1.4.2 User attribute definition (FIA_ATD.1)

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

- a) [user identifier(s):

- role;
 - *[identity;*
 - *password]*; and
- b) *[TOE Access Restriction(s):*
- *restricted source IP addresses; and*
 - *usage schedule]*.

5.1.4.3 Timing of authentication (for TOE Services) (FIA_UAU.1)

- FIA_UAU.1.1** The TSF shall allow [**ICMP Echo (ping), and ARP communications**] on behalf of the user to be performed before the user is authenticated.
- FIA_UAU.1.2** The TSF shall require each user to be successfully authenticated before allowing any other TSF-mediated actions on behalf of that user.

5.1.4.4 User authentication before any action (FIA_UAU.2)

- FIA_UAU.2.1** The TSF shall require the **Administrators and authorized IT entities** to be successfully authenticated before allowing any other TSF-mediated actions on behalf of these authorized users.

5.1.4.5 Multiple authentication mechanisms (FIA_UAU_(EXT).5)

- FIA_UAU_(EXT).5.1** The TSF shall provide a local authentication mechanism, *[and no other mechanism]* to perform user authentication.

5.1.4.6 User identification before any action (FIA_UID.2)

- 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.

5.1.4.7 User-subject binding (FIA_USB.1)

- FIA_USB.1.1** The TSF shall associate all user security attributes with subjects acting on the behalf of that authorized user.
- FIA_USB.1.2** The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users: [**security attributes are explicitly established based upon users security attributes**].
- FIA_USB.1.3** The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users: [**user security attributes cannot be changed within a session**].

5.1.5 Security management (FMT)

5.1.5.1 Management of security functions behavior (TSF non-cryptographic self-test) (FMT_MOF.1(1))

- FMT_MOF.1.1(1)** The TSF shall restrict the ability to determine and modify the behavior of the functions
- *[FPT_TST_(EXT).1]* to *[the security administrator]*.

5.1.5.2 Management of security functions behavior (cryptographic self-test) (FMT_MOF.1(2))

- FMT_MOF.1.1(2)** The TSF shall restrict the ability to enable, disable the functions

- [FPT_TST.1(1), and FPT_TST.1(2)]
to [the cryptographic administrator].

5.1.5.3 Management of security functions behavior (audit and alarms) (FMT_MOF.1(3))

- FMT_MOF.1.1(3)** The TSF shall restrict the ability to enable, disable, determine and modify the behavior of the functions
- [Security Audit (FAU_SAR)]
to [an administrator].

5.1.5.4 Management of security functions behavior (audit and alarms) (FMT_MOF.1(4))

- FMT_MOF.1.1(4)** The TSF shall restrict the ability to enable, disable, determine and modify the behavior of the functions
- [Security Audit Analysis (FAU_SAA); and
 - Security Audit (FAU_SEL)]
to [the security administrator].

5.1.5.5 Management of security functions behavior (audit and alarms) (FMT_MOF.1(5))

- FMT_MOF.1.1(5)** The TSF shall restrict the ability to enable, or disable the functions
- [Security Alarms (FAU_ARP)]
to [the security administrator].

5.1.5.6 Management of security functions behavior (available TOE-services for unauthenticated users) (FMT_MOF.1(6))

- FMT_MOF.1.1(6)** The TSF shall restrict the ability to enable, disable the functions [*ICMP*, [*none*]] to [the security administrator].

5.1.5.7 Management of security functions behavior (quota mechanism) (FMT_MOF.1(7))

- FMT_MOF.1.1(7)** The TSF shall restrict the ability to determine the behavior of the functions
- [Controlled connection-oriented resource allocation(FRU_RSA.1(2));
 - An administrator-specified network identifier;
 - set of administrator-specified network identifiers;
 - administrator-specified period of time]
- to [the security administrator].

5.1.5.8 Management of security attributes (FMT_MSA.1)

- FMT_MSA.1.1** The TSF shall enforce the [**TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, TRANSPARENT MODE FIREWALL SFP, ROUTE MODE FIREWALL SFP and UNAUTHENTICATED TOE SERVICES SFP**] to restrict the ability to [manipulate] the security attributes [referenced in the indicated policies] to [the security administrator].

5.1.5.9 Static attribute initialization (attributes/ruleset) (FMT_MSA.3(1))

- FMT_MSA.3.1(1)** The TSF shall enforce the [**TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, TRANSPARENT MODE FIREWALL SFP, and ROUTE MODE FIREWALL SFP**] to provide restrictive default values for the (security attributes) information flow policy ruleset that is (are) used to enforce the SFP.

FMT_MSA.3.2(1) The TSF shall allow the [Security Administrator] to specify alternative initial values to override the default values when an object or information is created.

5.1.5.10 Static attribute initialization (services) (FMT_MSA.3(2))

FMT_MSA.3.1(2) The TSF shall enforce the [UNAUTHENTICATED TOE SERVICES SFP] to provide restrictive default values for (security attributes) the set of TOE services available to unauthenticated users (that are used to enforce the SFP).

FMT_MSA.3.2(2) The TSF shall allow the [Security Administrator] to specify alternative initial values to override the default values when an object or information is created.

5.1.5.11 Management of TSF data (non-cryptographic, non-time TSF data) (FMT_MTD.1(1))

FMT_MTD.1.1(1) The TSF shall restrict the ability to [selection: change default, query, modify, delete, clear, *[none]*] all the [TSF data except cryptographic security data and the time and date used to form the time stamps in FPT_STM.1] to [the administrators or authorized IT entities].

5.1.5.12 Management of TSF data (cryptographic TSF data) (FMT_MTD.1(2))

FMT_MTD.1.1(2) The TSF shall restrict the ability to modify the [cryptographic security data] to [the cryptographic administrator].

5.1.5.13 Management of TSF data (time TSF data) (FMT_MTD.1(3))

FMT_MTD.1.1(3) The TSF shall restrict the ability to [set] the [time and date used to form the time stamps in FPT_STM.1] to [the security administrator or authorized IT entity].

5.1.5.14 Management of TSF data (Policy Rulesets) (FMT_MTD.1(4))

FMT_MTD.1.1(4) The TSF shall restrict the ability to [query, modify, delete, create] the [TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, TRANSPARENT MODE FIREWALL SFP, ROUTE MODE FIREWALL SFP and UNAUTHENTICATED TOE SERVICES SFP rules] to [the security administrator].

5.1.5.15 Management of limits on TSF data (transport-layer quotas) (FMT_MTD.2(1))

FMT_MTD.2.1(1) The TSF shall restrict the specification of the limits for [quotas on transport-layer connections] to [the security administrator].

FMT_MTD.2.2(1) The TSF shall take the following actions, if the TSF data are at, or exceed, the indicated limits: [**drops the packets and logs the event**].

5.1.5.16 Management of limits on TSF data (controlled connection-oriented quotas) (FMT_MTD.2(2))

FMT_MTD.2.1(2) The TSF shall restrict the specification of the limits for [quotas on controlled connection-oriented resources] to [the security administrator].

FMT_MTD.2.2(2) The TSF shall take the following actions, if the TSF data are at, or exceed, the indicated limits: [**log the connection attempt then allow or deny the connection based upon the configuration specified by the security administrator**].

5.1.5.17 Revocation (FMT_REV.1)

FMT_REV.1.1 The TSF shall restrict the ability to revoke security attributes associated with the [users, information flow policy ruleset, services available to unauthenticated users, **[none]**] within the TSC to [the security administrator].

FMT_REV.1.2 The TSF shall immediately enforce the

- a. [revocation of a user's role (Security Administrator, Cryptographic Administrator, Audit Administrator);
- b. changes to the information flow policy ruleset when applied;
- c. disabling of a service available to unauthenticated users;
- d. changes to the set of security associations with peer TOEs; and
- e. **[none]**].

5.1.5.18 Specification of Management Functions (FMT_SMF.1)

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

1. restrict the ability to invoke determine and modify the behavior of functions: [the TSF Self-Test (FPT_TST_(EXT).1)] to [the Security Administrator];
2. restrict the ability to enable, disable the functions TSF Self-Test (Crypto Self-Test) (FPT_TST.1(1), and) Key Generation Self-Test (FPT_TST.1(2)) to the Cryptographic Administrator;
3. restrict the ability to enable, disable, determine and modify the behavior of the functions Security Audit (FAU_SAR) to an Administrator;
4. restrict the ability to enable, disable, determine and modify the behavior of the functions Security Audit Analysis (FAU_SAA); and Security Audit (FAU_SEL) to the Security Administrator;
5. restrict the ability to enable, or disable the functions Security Alarms (FAU_ARP) to the Security Administrator;
6. restrict the ability to enable, disable the functions **[ICMP]** to [the Security Administrator];
7. restrict the ability to determine the behavior of the functions An administrator-specified network identifier; set of administrator-specified network identifiers; administrator-specified period of time] to [the Security Administrator];
8. enforce the [**TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, TRANSPARENT MODE FIREWALL SFP, ROUTE MODE FIREWALL SFP and UNAUTHENTICATED TOE SERVICES SFP**] to restrict the ability to manipulate the security attributes referenced in the indicated policies to an Administrator;
9. enforce the [**TRANSPARENT MODE VPN SFP and ROUTE MODE VPN SFP**] to provide restrictive default values for the information flow policy rule set security attributes that is used to enforce the SFP;
10. enforce the [**UNAUTHENTICATED TOE SERVICES SFP**] to provide restrictive default values security attributes that are used to enforce the SFP;
11. restrict the ability to [change default, query, modify, delete, clear, **[none]**] all the [TSF data except cryptographic security data and the time and date used to form the time stamps in FPT_STM.1] to [the administrators or authorized IT entities];
12. restrict the ability to modify the cryptographic security data to the Cryptographic Administrator;
13. restrict the ability to set the time and date used to form the time stamps in [FPT_STM.1] to the Security Administrator or authorized IT entity;
14. restrict the ability to query, modify, delete, create, [none] the VPN Policy rules to the Security Administrator;
15. restrict the specification of the limits for quotas on transport-layer connections to the Security Administrator;
16. restrict the specification of the limits for quotas on controlled connection-oriented resources to the Security Administrator;

17. [enforce the **TRANSPARENT MODE FIREWALL SFP and ROUTE MODE FIREWALL SFP to provide restrictive default values for the information flow policy rule set security attributes that is used to enforce the SFP**].

5.1.5.19 Restrictions on security roles (FMT_SMR.2)

FMT_SMR.2.1 The TSF shall maintain the roles: [

- Security Administrator;
- Cryptographic Administrator (i.e., users authorized to perform cryptographic initialization and management functions);
- Audit Administrators;
- authorized IT entities; and
- *[none]*.

FMT_SMR.2.2 The TSF shall be able to associate users with roles.

FMT_SMR.2.3 The TSF shall ensure that the conditions [

1. All roles shall be able to administer the TOE locally;
2. all roles shall be able to administer the TOE remotely;
3. all roles are distinct; that is, there shall be no overlap of operations performed by each role, with the following exceptions:
 - all administrators can review the audit trail; and
 - all administrators can invoke the self-tests

] are satisfied.

5.1.6 Protection of the TSF (FPT)

5.1.6.1 Manual recovery (FPT_RCV.1)

FPT_RCV.1.1 After a [failure or service discontinuity], the TSF shall enter a maintenance mode where the ability to return the TOE to a secure state is provided.

5.1.6.2 Replay detection (FPT_RPL.1)

FPT_RPL.1.1 The TSF shall detect replay for the following entities: [TSF data and security attributes].

FPT_RPL.1.2 The TSF shall perform [reject data, audit event and **[no other actions]**] when replay is detected.

5.1.6.3 Reliable time stamps (FPT_STM.1)

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

5.1.6.4 TSF testing (FPT_TST_(EXT).1)

FPT_TST_(EXT).1.1 The TSF shall run a suite of self-tests during initial start-up and also either periodically during normal operation, or at the request of an authorized administrator to demonstrate the correct operation of the TSF.

FPT_TST_(EXT).1.2 The TSF shall provide authorized administrators with the capability to verify the integrity of stored TSF executable code through the use of the TSF-provided cryptographic services.

5.1.6.5 TSF testing (for cryptography) (FPT_TST.1(1))

FPT_TST.1.1(1) The TSF shall run the suite of self-tests in accordance with FIPS PUB 140-2 and ~~Appendix F~~ **Section 10** of this ~~profile~~ **Security Target** during initial start-up (on power on), at the request of the cryptographic administrator (on demand), under various conditions defined in section

4.9.1 of FIPS 140-2, and periodically (at least once a day) to demonstrate the correct operation of the following cryptographic functions:

- a) key error detection;
- b) cryptographic algorithms;
- c) RNG/PRNG.

FPT_TST.1.2(1) The TSF shall provide authorized cryptographic administrators with the capability to verify the integrity of TSF data related to the cryptography by using TSF-provided cryptographic functions.

FPT_TST.1.3(1) The TSF shall provide authorized cryptographic administrators with the capability to verify the integrity of stored TSF executable code related to the cryptography by using TSF-provided cryptographic functions.

5.1.6.6 TSF testing (for key generation components) (FPT_TST.1(2))

FPT_TST.1.1(2) The TSF shall perform self tests immediately after generation of a key to demonstrate the correct operation of each key generation component. If any of these tests fails, that generated key shall not be used, the cryptographic module shall react as required by FIPS PUB 140-2 for failing a self-test, and this event will be audited.

FPT_TST.1.2(2) The TSF shall provide authorized cryptographic administrators with the capability to verify the integrity of TSF data related to the key generation by using TSF-provided cryptographic functions.

FPT_TST.1.3(2) The TSF shall provide authorized cryptographic administrators with the capability to verify the integrity of stored TSF executable code related to the key generation by using TSF-provided cryptographic functions.

5.1.7 Resource utilization (FRU)

5.1.7.1 Maximum quotas (transport-layer quotas) (FRU_RSA.1(1))

FRU_RSA.1.1(1) The TSF shall enforce maximum quotas of the following resources: [transport-layer representation] that **a source subject identifier users** can use over a specified period of time.

5.1.7.2 Maximum quotas (controlled connection-oriented quotas) (FRU_RSA.1(2))

FRU_RSA.1.1(2) The TSF shall enforce administrator-specified maximum quotas of the following resources: [*transport-layer sessions traversing the firewall*] that users associated with an administrator-specified network identifier and a set of administrator-specified network identifiers can use over an administrator-specified period of time.

5.1.8 TOE access (FTA)

5.1.8.1 TSF-initiated session locking (FTA_SSL.1)

FTA_SSL.1.1 The TSF shall lock a local interactive session after [a Security Administrator-specified time period of inactivity] by:

- a) clearing or overwriting display devices, making the current contents unreadable;
- b) disabling any activity of the user's data access/display devices other than unlocking the session.

FTA_SSL.1.2 The TSF shall require the following events to occur prior to unlocking the session: [the administrator to re-authenticate].

5.1.8.2 User-initiated locking (FTA_SSL.2)

FTA_SSL.2.1 The TSF shall allow user-initiated locking of the user's own local interactive session, by:

- clearing or overwriting display devices, making the current contents unreadable;
- disabling any activity of the user's data access/display devices other than unlocking the session.

FTA_SSL.2.2 The TSF shall require the following events to occur prior to unlocking the session [the administrator to re-authenticate].

5.1.8.3 TSF-initiated termination (FTA_SSL.3)

FTA_SSL.3.1 The TSF shall terminate a remote session after a [Security Administrator-configurable time interval of session inactivity].

5.1.8.4 Default TOE access banners (FTA_TAB.1)

FTA_TAB.1.1 Before establishing an administrator session the TSF shall display only a Security Administrator-specified advisory notice and consent warning message regarding unauthorized use of the TOE.

5.1.8.5 TOE session establishment (FTA_TSE.1)

FTA_TSE.1.1 The TSF shall be able to deny establishment of an administrator session based on [location, time, and day].

5.1.9 Trusted path/channels (FTP)

5.1.9.1 Inter-TSF trusted channel (Prevention of Disclosure) (FTP_ITC.1(1))

FTP_ITC.1.1(1) The TSF shall use encryption to provide a trusted communication channel between itself and authorized IT entities that is logically distinct from other communication channels and provides assured identification of its end points and protection of the channel data from disclosure.

FTP_ITC.1.2(1) The TSF shall permit the TSF, or the authorized IT entities to initiate communication via the trusted channel.

FTP_ITC.1.3(1) The TSF shall initiate communication via the trusted channel for [all authentication functions, *[none]*].

5.1.9.2 Inter-TSF trusted channel (Prevention of Modification) (FTP_ITC.1(2))

FTP_ITC.1.1(2) The TSF shall use a cryptographic signature to provide a trusted communication channel between itself and authorized IT entities that is logically distinct from other communication channels and provides assured identification of its end points and detection of the modification of data.

FTP_ITC.1.2(2) The TSF shall permit the TSF, or the authorized IT entities to initiate communication via the trusted channel.

FTP_ITC.1.3(2) The TSF shall initiate communication via the trusted channel for [all authentication functions, *[none]*].

5.1.9.3 Trusted Path (Prevention of Disclosure) (FTP_TRP.1(1))

FTP_TRP.1.1(1) The TSF shall provide an encrypted communication path between itself and remote administrators that is logically distinct from other communication paths and provides assured identification of its end points and protection of the communicated data from disclosure.

FTP_TRP.1.2(1) The TSF shall permit administrators to initiate communication via the trusted path.

FTP_TRP.1.3(1) The TSF shall require the use of the trusted path for all remote administration actions.

5.1.9.4 Trusted Path (Detection of Modification) (FTP_TRP.1(2))

FTP_TRP.1.1(2) The TSF shall use a cryptographic signature to provide a communication path between itself and administrators that is logically distinct from other communication paths and provides assured identification of its end points and detection of the modification of data.

FTP_TRP.1.2(2) The TSF shall permit administrators to initiate communication via the trusted path.

FTP_TRP.1.3(2) The TSF shall require the use of the trusted path for all remote administration actions.

5.2 TOE Security Assurance Requirements

The security assurance requirements for the TOE are the EAL 4 augmented with ADV_FSP.5, ADV_INT.3, ADV_TDS.4, ALC_FLR.2, ATE_DPT.3, and AVA_VAN.3 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.5: Complete semi-formal functional specification with additional error information
	ADV_IMP.1: Implementation representation of the TSF
	ADV_INT.3: Minimally complex internals
	ADV_TDS.4: Semiformal modular design
AGD: Guidance documents	AGD_OPE.1: Operational user guidance
	AGD_PRE.1: Preparative procedures
ALC: Life-cycle support	ALC_CMC.4: Production support, acceptance procedures and automation
	ALC_CMS.4: Problem tracking CM coverage
	ALC_DEL.1: Delivery procedures
	ALC_DVS.1: Identification of security measures
	ALC_FLR.2: Flaw reporting procedures
	ALC_LCD.1: Developer defined life-cycle model
	ALC_TAT.1: Well-defined development tools
ATE: Tests	ATE_COV.2: Analysis of coverage
	ATE_DPT.3: Testing: modular design
	ATE_FUN.1: Functional testing
	ATE_IND.2: Independent testing - sample
AVA: Vulnerability assessment	
	AVA_VAN.3: Focused vulnerability analysis

Table 2 Medium 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 Complete semi-formal functional specification with additional error information (ADV_FSP.5)

- ADV_FSP.5.1d** The developer shall provide a functional specification.
- ADV_FSP.5.2d** The developer shall provide a tracing from the functional specification to the SFRs.
- ADV_FSP.5.1c** The functional specification shall completely represent the TSF.
- ADV_FSP.5.2c** The functional specification shall describe the TSFI using a semi-formal style.
- ADV_FSP.5.3c** The functional specification shall describe the purpose and method of use for all TSFI.
- ADV_FSP.5.4c** The functional specification shall identify and describe all parameters associated with each TSFI.
- ADV_FSP.5.5c** The functional specification shall describe all actions associated with each TSFI.
- ADV_FSP.5.6c** The functional specification shall describe all direct error messages that may result from an invocation of each TSFI.
- ADV_FSP.5.7c** The functional specification shall describe all error messages that do not result from an invocation of a TSFI.
- ADV_FSP.5.8c** The functional specification shall provide a rationale for each error message contained in the TSF implementation yet does not result from an invocation of a TSFI.
- ADV_FSP.5.9c** The tracing shall demonstrate that the SFRs trace to TSFIs in the functional specification.
- ADV_FSP.5.1e** The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
- ADV_FSP.5.2e** The evaluator shall determine that the functional specification is an accurate and complete instantiation of the SFRs.

5.2.1.3 Implementation representation of the TSF (ADV_IMP.1)

- ADV_IMP.1.1d** The developer shall make available the implementation representation for the entire TSF.
- ADV_IMP.1.2d** The developer shall provide a mapping between the TOE design description and the sample of the implementation representation.
- ADV_IMP.1.1c** The implementation representation shall define the TSF to a level of detail such that the TSF can be generated without further design decisions.
- ADV_IMP.1.2c** The implementation representation shall be in the form used by the development personnel.
- ADV_IMP.1.3c** The mapping between the TOE design description and the sample of the implementation representation shall demonstrate their correspondence.
- ADV_IMP.1.1e** The evaluator shall confirm that, for the selected sample of the implementation representation, the information provided meets all requirements for content and presentation of evidence.

5.2.1.4 Minimally complex internals (ADV_INT.3)

- ADV_INT.3.1d** The developer shall design and implement the entire TSF such that it has well-structured internals.
- ADV_INT.3.2d** The developer shall provide an internals description and justification.
- ADV_INT.3.1c** The justification shall explain the characteristics used to judge the meaning of 'well-structured' and 'complex'.
- ADV_INT.3.2c** The TSF internals description shall demonstrate that the entire TSF is well-structured and is not overly complex.

- ADV_INT.3.1e** The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
- ADV_INT.3.2e** The evaluator shall perform an internals analysis on the entire TSF.

5.2.1.5 Semiformal modular design (ADV_TDS.4)

- ADV_TDS.4.1d** The developer shall provide the design of the TOE.
- ADV_TDS.4.2d** 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.4.1c** The design shall describe the structure of the TOE in terms of subsystems.
- ADV_TDS.4.2c** The design shall describe the TSF in terms of modules, designating each module as SFR-enforcing, SFR-supporting, or SFR-non-interfering.
- ADV_TDS.4.3c** The design shall identify all subsystems of the TSF.
- ADV_TDS.4.4c** The design shall provide a semiformal description of each subsystem of the TSF, supported by informal, explanatory text where appropriate.
- ADV_TDS.4.5c** The design shall provide a description of the interactions among all subsystems of the TSF.
- ADV_TDS.4.6c** The design shall provide a mapping from the subsystems of the TSF to the modules of the TSF.
- ADV_TDS.4.7c** The design shall describe each SFR-enforcing and SFR-supporting module in terms of its purpose and interaction with other modules.
- ADV_TDS.4.8c** The design shall describe each SFR-enforcing and SFR-supporting module in terms of its SFR-related interfaces, return values from those interfaces, interaction with and called interfaces to other modules.
- ADV_TDS.4.9c** The design shall describe each SFR-non-interfering module in terms of its purpose and interaction with other modules.
- ADV_TDS.4.10c** The mapping shall demonstrate that all behaviour described in the TOE design is mapped to the TSFIs that invoke it.
- ADV_TDS.4.1e** The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
- ADV_TDS.4.2e** 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 Production support, acceptance procedures and automation (ALC_CMC.4)

ALC_CMC.4.1d The developer shall provide the TOE and a reference for the TOE.

ALC_CMC.4.2d The developer shall provide the CM documentation.

ALC_CMC.4.3d The developer shall use a CM system.

ALC_CMC.4.1c The TOE shall be labeled with its unique reference.

ALC_CMC.4.2c The CM documentation shall describe the method used to uniquely identify the configuration items.

ALC_CMC.4.3c The CM system shall uniquely identify all configuration items.

ALC_CMC.4.4c The CM system shall provide automated measures such that only authorised changes are made to the configuration items.

ALC_CMC.4.5c The CM system shall support the production of the TOE by automated means.

ALC_CMC.4.6c The CM documentation shall include a CM plan.

ALC_CMC.4.7c The CM plan shall describe how the CM system is used for the development of the TOE.

ALC_CMC.4.8c The CM plan shall describe the procedures used to accept modified or newly created configuration items as part of the TOE.

ALC_CMC.4.9c The evidence shall demonstrate that all configuration items are being maintained under the CM system.

ALC_CMC.4.10c The evidence shall demonstrate that the CM system is being operated in accordance with the CM plan.

ALC_CMC.4.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.3.2 Problem tracking CM coverage (ALC_CMS.4)

ALC_CMS.4.1d The developer shall provide a configuration list for the TOE.

ALC_CMS.4.1c The configuration list shall include the following: the TOE itself; the evaluation evidence required by the SARs; the parts that comprise the TOE; the implementation representation; and security flaw reports and resolution status.

ALC_CMS.4.2c The configuration list shall uniquely identify the configuration items.

ALC_CMS.4.3c For each TSF relevant configuration item, the configuration list shall indicate the developer of the item.

ALC_CMS.4.1e 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 Identification of security measures (ALC_DVS.1)

- ALC_DVS.1.1d** The developer shall produce development security documentation.
- ALC_DVS.1.1c** The development security documentation shall describe all the physical, procedural, personnel, and other security measures that are necessary to protect the confidentiality and integrity of the TOE design and implementation in its development environment.
- ALC_DVS.1.1e** The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
- ALC_DVS.1.2e** The evaluator shall confirm that the security measures are being applied.

5.2.3.5 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.

5.2.3.6 Developer defined life-cycle model (ALC_LCD.1)

- ALC_LCD.1.1d** The developer shall establish a life-cycle model to be used in the development and maintenance of the TOE.
- ALC_LCD.1.2d** The developer shall provide life-cycle definition documentation.
- ALC_LCD.1.1c** The life-cycle definition documentation shall describe the model used to develop and maintain the TOE.
- ALC_LCD.1.2c** The life-cycle model shall provide for the necessary control over the development and maintenance of the TOE.
- ALC_LCD.1.1e** The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.3.7 Well-defined development tools (ALC_TAT.1)

- ALC_TAT.1.1d** The developer shall identify each development tool being used for the TOE.
- ALC_TAT.1.2d** The developer shall document the selected implementation-dependent options of each development tool.
- ALC_TAT.1.1c** Each development tool used for implementation shall be well-defined.
- ALC_TAT.1.2c** The documentation of each development tool shall unambiguously define the meaning of all statements as well as all conventions and directives used in the implementation.
- ALC_TAT.1.3c** The documentation of each development tool shall unambiguously define the meaning of all implementation-dependent options.

ALC_TAT.1.1e 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 Analysis of coverage (ATE_COV.2)

ATE_COV.2.1d The developer shall provide an analysis of the test coverage.

ATE_COV.2.1c The analysis of the test coverage shall demonstrate the correspondence between the tests in the test documentation and the TSFIs in the functional specification.

ATE_COV.2.2c The analysis of the test coverage shall demonstrate that all TSFIs in the functional specification have been tested.

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

5.2.4.2 Testing: modular design (ATE_DPT.3)

ATE_DPT.3.1d The developer shall provide the analysis of the depth of testing.

ATE_DPT.3.1c The analysis of the depth of testing shall demonstrate the correspondence between the tests in the test documentation and the TSF subsystems and modules in the TOE design.

ATE_DPT.3.2c The analysis of the depth of testing shall demonstrate that all TSF subsystems in the TOE design have been tested.

ATE_DPT.3.3c The analysis of the depth of testing shall demonstrate that all modules in the TOE design have been tested.

ATE_DPT.3.1e The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.

5.2.4.3 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.

5.2.4.4 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 Methodical vulnerability analysis (AVA_VAN.3)

AVA_VAN.3.1d The developer shall provide the TOE for testing.

- AVA_VAN.3.1c** The TOE shall be suitable for testing.
- AVA_VAN.3.1e** The evaluator shall confirm that the information provided meets all requirements for content and presentation of evidence.
- AVA_VAN.3.2e** The evaluator shall perform a search of public domain sources to identify potential vulnerabilities in the TOE.
- AVA_VAN.3.3e** The evaluator shall perform an independent vulnerability analysis of the TOE using the guidance documentation, functional specification, TOE design, security architecture description and implementation representation to identify potential vulnerabilities in the TOE.
- AVA_VAN.3.4e** 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 Enhanced-Basic attack potential.

6. TOE Summary Specification

This chapter describes the security functions.

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 or 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.

Buffer storage on the device is broken into the following categories. There are two buffers for event logs, one for basic logs and one for alarms. There are also two buffers for traffic & 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 also 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.

Security appliances provide memory to hold a fixed maximum number of audit records and then once the storage limit is reached, the audit mechanism 'wraps' or acts as a first-in-first-out (FIFO) stack, when overwriting the oldest audit information in the storage device with the new audit information. 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. When so configured, security appliances do follow every write to an audit log 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, so that all writes to the backup device will be serviced without audit data loss.

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, average traffic flow, and lulls where traffic that causes audited events are low. The worst case for audit loss would occur if memory wrote an audit record in the last available location, and a burst of audited events occurred before they 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 memory can never lose audit information in that no audit records can ever be "dropped" or not written. Additionally, the security appliances can be configured to notify the administrator (via an 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. 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 will be 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. This feature is available as a configuration option.

The Security, Audit and Cryptographic administrators have access to the audit logs and memory 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 include:

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 exceeded (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 Security Administrator may also define an alarm threshold for monitoring events on the system. Once a threshold is exceeded, an alarm is generated. The following events can be assigned an alarm threshold:

- a) Security Administrator specified number of authentication failures
- b) Security Administrator specified number of Information Flow policy violations by an individual presumed source network identifier (e.g., IP address) within an administrator specified time period
- c) Security Administrator specified number of Information Flow policy violations to an individual destination network identifier within an administrator specified time period
- d) Security Administrator specified number of Information Flow policy violations to an individual destination subject service identifier (e.g., TCP port) within an administrator specified time period
- e) Security Administrator specified Information Flow policy rule, or group of rule violations within an administrator specified time period
- f) Any detected replay of TSF data or security attributes
- g) Any failure of the cryptomodule self-tests (FPT_TST.1(1))
- h) Any failure of the other key generation self-tests (FPT_TST.1(2))
- i) Any failure of the other TSF self-tests (FPT_TST_(EXT). 1)
- j) Security Administrator specified number of encryption failures
- k) Security Administrator specified number of decryption failures
- l) Security Administrator specified number of Phase 1 authentication failures when negotiating the Internet Key Exchange protocol, and
- m) Security Administrator specified number of failures occurring during Phase 2 negotiation.

When an alarm is generated, a message is displayed at the local console, at remote administrator sessions that currently exist and at remote administrator sessions that are initiated before the alarm has been acknowledged.

The alarm message contains an alarm id. To retrieve the details of the alarm and acknowledge it, the administrator executes a command in which the alarm id is specified. If configured, alarms are accompanied by a terminal chime that repeats until all alarms have been acknowledged. To ensure that the alarm id remains on screen, it is reprinted with every command line prompt.

When an alarm is acknowledged, an acknowledgement notice is printed on the local console and any remote consoles that received the alarm, listing the alarm id, time of acknowledgement and user performing the

acknowledgement. The TOE provides a Command Line Interface (CLI) for administrators to review the logs that record audited events, using the CLI 'get' commands. The logs display the date, time, level, and description for each event.

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:

- a) user identity
- b) event type
- c) network identifier
- d) subject service identifier
- e) success of auditable security events
- f) failure of auditable security events, and
- g) 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_ARP.1: Potential security violations can cause alarms that are displayed at the local console, at remote administrator sessions that currently exist and at remote administrator sessions that are initiated before the alarm has been acknowledged. These alarms can be configured to be audible.
- FAU_ARP_ACK_(EXT).1: Alarm messages sent to an administrator are displayed at the local console, at remote administrator sessions that currently exist and at remote administrator sessions that are initiated before the alarm has been acknowledged. These alarm messages provide information about the potential security violation that caused the alarm message (including the pertinent audit record). When an alarm is acknowledged, an acknowledgement notice is printed. If configured, audible alarms are accompanied by a terminal chime that repeats until all alarms have been acknowledged. To ensure that the alarm id remains on screen, it is reprinted with every command line prompt.
- FAU_GEN.1-NIAP-0410: The TSF generates audit records for the events listed above.
- FAU_GEN.2-NIAP-0410: 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 in Section 9 of this security target.
- FAU_SAA.1-NIAP-0407: The TSF can establish alarm thresholds for the events listed above. Once the number of events indicated by the threshold has been exceeded, an alarm is generated.
- FAU_SAR.1: Administrators can read audit data using command line operations.
- FAU_SAR.2: Access to the audit logs is available to the Security, Audit and Cryptographic administrators, and only to those administrators.
- 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_SEL.1-NIAP-0407: The 'set log' command allows the security administrator to exclude certain events from being logged based on the specific attributes of the audited event. The selectable attributes are listed above.
- FAU_STG.1-NIAP-0429: 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.3: The security appliances can be configured to notify the administrator when the amount of local audit storage used has reached a specified percentage. The notification takes the form of an alarm message as required by FAU_ARP.1.
- FAU_STG.NIAP-0414-1-NIAP-0429: The TSF allows the security administrator to configure the system such that the TOE behaves in either of the following manners:
 - 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.
 - In the event that the TOE is configured to overwrite the oldest records when local audit storage is consumed, the TSF overwrites the oldest stored audit record.

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.

The FIPS-approved cryptomodule 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 TSF supports the manual and electronic distribution of cryptographic keys. Support for distribution of symmetric keys is in accordance with FIPS PUB 171 (Key Management Using ANSI X9.17). In ScreenOS private asymmetric keys are only distributed using manual distribution methods. ScreenOS automatically distributes public asymmetric key material (certificates and/or keys) in accordance with NSA-certified DoD PKI for public key distribution using NSA-approved certificate schemes that meet the 'PKI Roadmap for the DoD' and 'DoD X.509 Certificate Policy'.

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

- FCS_BCM_(EXT).1: ScreenOS is implemented as a combination of software and hardware. Cryptographic processing occurs only in software on some platforms, while others use commercial or custom ASIC's to accelerate cryptographic operations. ScreenOS meets FIPS 140-2 requirements by allowing the administrator to enable a FIPS operating mode. When the administrator enables FIPS mode, the device zeroes itself and reboots. When operating in FIPS mode, ScreenOS offers a restricted set of configuration options and enforces a more restrictive set of security practices.

All security appliances comprising the TOE are FIPS 140-2 Security Level 2 validated, with Security Level 3 in the following areas:

- Cryptographic Module Ports and Interfaces
- Roles, Services and Authentication
- Cryptographic Key Management
- Design Assurance

The certificate numbers for the completed FIPS evaluation are:

- SSG 5 and SSG 20, #1170
 - SSG 140, #1171
 - SSG 320M and SSG 350M, #1172
 - SSG 520M and SSG 550M, #1173
 - NetScreen-ISG1000 and NetScreen-ISG2000, #1169
 - NetScreen-5200 and NetScreen-5400, #1168
- FCS_CKM.1(1): The TOE allows VPN connections to be configured in one of three ways:
 1. Manual keys
 2. Auto-key IKE with pre-shared keys for authentication
 3. Auto-key IKE with certificates for authentication

Manually keyed VPNs do not draw upon the RNG; the administrator explicitly configures the key values. Auto-key IKE VPNs draw upon the RNG during the key agreement process. The TOE employs a FIPS 140-2 validated software RNG.

Auto-IKE VPNs perform Diffie-Hellman key agreement to generate symmetric keys. The TOE supports only Diffie-Hellman key agreement. The TOE only generates symmetric keys in the context of a Diffie-Hellman key agreement. Generated keys are protected in a manner compliant with SP 800-57, section 6.1.

- FCS_CKM.1(2): The TSF generates 2048-bit DH keys, which are equivalent in strength to 112 bit symmetric keys and 256-bit ECDSA keys, which are equivalent or greater in strength to 128 bit symmetric keys. Prime numbers are generated using a method that complies with ANSI X9.80 and X9.42. The FIPS-approved ECDSA algorithm is defined by ANSI X9.62-1998. Support for distribution of symmetric keys is in accordance with NIST SP 800-57.
- FCS_CKM.2: The TSF supports the manual and electronic distribution of cryptographic keys. Support for distribution of symmetric keys is in accordance with NIST SP 800-57 and 800-56A.
- FCS_CKM_(EXT).2: The TSF performs key input and output in accordance with FIPS 140-2, level 3. Keys are associated with the correct entity through means such as a Security Parameters Index (SPI), a fully qualified domain name (FQDN) or a connection index. A parity check is performed whenever a key is internally transferred. All keys are encrypted when not in use. Administrators can define a period of inactivity, after which the ScreenOS will destroy non-persistent cryptographic keys. When no longer needed, memory space used by a key is overwritten by a pseudo-random bit pattern. ScreenOS does not provide a mechanism to archive expired private signature keys.
- FCS_CKM.4: The TSF supports a zeroization command line option that destroys all keying material, overwriting it with a pseudo-random bit pattern. In addition, this command resets the device to the factory default configuration and restarts it. Further, freed key storage memory is always zeroized whenever the key is moved, copied or deleted.
- 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): ScreenOS supports the following digital signature algorithms:
 - ECDSA with a key size of 256 bits using the NIST curve, P-256.
- FCS_COP.1(3): ScreenOS supports cryptographic hashing via the SHA-256 algorithm.
- FCS_COP.1(4): ScreenOS supports cryptographic key agreement via Diffie-Hellman groups 1, 2, 5 and 14. These Diffie-Hellman groups provide a key agreement algorithm for finite field-based key agreement algorithm and cryptographic key sizes (modulus) of 2048 bits, that meet ANSI X9.42-2001.

- FCS_COP_(EXT).1: ScreenOS performs random number generation via the ANSI X9.31 algorithm, using multiple software-generated seed values combined with SHA-1 hashing. The random number generator meets testing requirements from FIPS PUB 140-2 and 180-2; NIST Special Publication 800-22; and TSF self tests from the Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile.
- FCS_IKE_(EXT).1: ScreenOS supports IKE v1 in Main and Aggressive modes as per RFC 2409. Authentication is available through pre-shared key, RSA, DSA and ECDSA certificates. All random values are generated using the FIPS-approved RNG.

ScreenOS computes the value of SKEYID (as defined in RFC 2409), using SHA-1 as the pseudo-random function. ScreenOS authenticates using the methods for

- Signatures: $SKEYID = sha(Ni_b | Nr_b, g^{xy})$
- Pre-shared keys: $SKEYID = sha(\text{pre-shared-key}, Ni_b | Nr_b)$
- Authentication using Public key encryption, computing SKEYID as follows: $= sha(sha(Ni_b | Nr_b), CKY-I | CKY-R)$

ScreenOS computes authenticated keying material as follows:

- $SKEYID_d = sha(SKEYID, g^{xy} | CKY-I | CKY-R | 0);$
- $SKEYID_a = sha(SKEYID, SKEYID_d | g^{xy} | CKY-I | CKY-R | 1);$ and
- $SKEYID_e = sha(SKEYID, SKEYID_a | g^{xy} | CKY-I | CKY-R | 2).$ To authenticate the Phase 1 exchange, ScreenOS generates HASH-I if it is the initiator and HASH-R if it is the responder, according to RFC 2409 and as defined by the SFR in this security target.

ScreenOS authenticates IKE Phase 1 using authentication with digital signatures or with a pre-shared key as defined by the SFR in this security target. For digital signatures ScreenOS can apply an RSA signature to HASH-I or HASH-R if the signature is PKCS#1 encoded as defined by the SFR in this security target. ScreenOS can also use X.509 Version 3 certificates.

ScreenOS computes hashes in the following way

$$HASH(1) = sha(SKEYID_a, M-ID | [\text{any ISAKMP payload after HASH(1) header contained in the message}])$$

$$HASH(2) = sha(SKEYID_a, M-ID | Ni_b | [\text{any ISAKMP payload after HASH(2) header contained in the message}])$$

$$HASH(3) = sha(SKEYID_a, 0 | M-ID | Ni_b | Nr_b)$$

ScreenOS computes keying material during quick mode using perfect forward secrecy. ScreenOS supports the following ID types: ID_IPV4_ADDR, ID_FQDN, ID_USER_FQDN, ID_IPV4_ADDR_SUBNET, ID_IPV6_ADDR, ID_IPV6_ADDR_SUBNET. ScreenOS provides cryptographic key establishment techniques in accordance with RFC 2409 as follows(s):

- Phase 1, the establishment of a secure authenticated channel between the TOE and another remote VPN endpoint, is performed using one of the following, as configured by the security administrator:
 - Main Mode
 - Aggressive Mode

New Group mode shall include the private group 14, 2048-bit MOD P for the Diffie-Hellman key exchange.

- Phase 2, negotiation of security services for IPsec, is done using Quick Mode, using SHA-1 as the pseudo-random function. Quick Mode generates key material that provides perfect forward secrecy. ScreenOS requires the nonce, and the x of g^{xy} be randomly generated using FIPS-approved random number generator when computation is being performed. The minimum size of x is twice the number of bits of the strength level associated with the negotiated DH group per table 2 of NIST SP 800-57.

The nonce sizes are 256 bits. Nonces are generated in a manner such that the probability that a specific nonce value will be repeated during the life of a specific IPsec SA is less than 1 in $2^{(bit\ strength\ of\ the\ negotiated\ DH\ group)}$.

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 TRANSPARENT MODE VPN SFP and ROUTE MODE 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 Gateway² (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. The TRANSPARENT MODE VPN SFP and ROUTE MODE 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:

² The RSH ALG filtering is not supported when used with port address translation.

- 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.

- 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 TRANSPARENT MODE FIREWALL SFP and ROUTE MODE FIREWALL 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 Gateway³ (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. The TRANSPARENT MODE FIREWALL SFP only supports the actions to permit, deny or reject. The ROUTE MODE FIREWALL SFP 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:

³ The RSH ALG filtering is not supported when used with port 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;
- 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.

6.1.3.3 TOE Services

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 (telnet, ping, snmp, HTTP) 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.

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

- FDP_IFC.1(1): The TRANSPARENT MODE VPN SFP applies to traffic to or from a network interface configured in Transparent mode that is using a VPN tunnel.
- FDP_IFC.1(2): The ROUTE MODE VPN SFP applies to traffic to or from a network interface configured in Route or NAT mode that is using a VPN tunnel.
- FDP_IFC.1(3): The UNAUTHENTICATED TOE SERVICES SFP applies to traffic directed at the TOE.
- FDP_IFC.1(4): The TRANSPARENT MODE FIREWALL SFP applies to traffic to or from a network interface configured in Transparent mode that is not using a VPN tunnel and is not directed at the TOE.
- FDP_IFC.1(5): The ROUTE MODE FIREWALL SFP applies to traffic to or from a network interface configured in Route or NAT mode that is not using a VPN tunnel and is not directed at the TOE.
- FDP_IFF.1(1): The TRANSPARENT MODE 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 TRANSPARENT MODE VPN SFP are mapped to actions in the corresponding SFRs as follows:
 - the pass action maps to permit and nat;
 - 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_1FF.1(2): The ROUTE MODE 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 ROUTE MODE 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.

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

- FDP_1FF.1(3): The UNAUTHENTICATED TOE SERVICES SFP is enforced on information flows directed at the TOE. The TSF enforces an 'access policy' on information flows directed at the TOE that either accepts or rejects network packets. Security appliances support multiple policies based upon a 'zone', using the rules defined in the 'VPN Policies' section above. The TOE supports only four types of direct inbound communication pathways. These inbound communication pathways are unauthenticated ICMP echo, unauthenticated ARP communications, SSH protected remote administration sessions and VPN connections to authorized IT entities. These SSH and VPN pathways are enforced using the algorithm described in the 'TOE Services' 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. Policy-Based Address Translation is applied to IPv4 interfaces configured for NAT.
- FDP_1FF.1(4): The TRANSPARENT MODE FIREWALL 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.

The actions defined in an 'access policy' that enforces the TRANSPARENT MODE FIREWALL SFP are mapped to actions in the corresponding SFRs as follows:

- the pass action maps to permit;
- the drop action maps to both deny and reject.
- FDP_1FF.1(5): The ROUTE MODE FIREWALL 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 '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.

The actions defined in an 'access policy' that enforces the ROUTE MODE FIREWALL SFP are mapped to actions in the corresponding SFRs as follows:

- the pass action maps to permit and nat;

- the drop action maps to both deny and reject.

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

- 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.

Key material resources are distributed and managed using the security appliances' IPSec capabilities. All temporary storage associated with key material is handled in the same manner since it is encapsulated within packets. 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 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 utilize encryption and digital signatures to 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. An administrator can also be restricted to only login from specified IP addresses or according to a usage schedule.
- 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 external servers, only ICMP echo and ARP communications are unauthenticated.
- FIA_UAU.2: The first and only interface presented to an administrator is a command line requesting a user identifier and password. There are no other interfaces to the TOE presented to a user.
- FIA_UAU_(EXT).5: 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.

- FIA_UID.2: Users of the TOE include administrators and external servers. Administrators must present a valid user id during the login process. External servers are identified by their source IP address, which is provided as part of every IP packet.
- FIA_USB.1: 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. Security attributes are explicitly established based upon user's security attributes upon login and cannot be changed within a session.

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 Security, Cryptographic and Audit administrative roles. Once the three accounts have been created, there is no overlap between the privileges available to them, except the ability to review the audit trail and invoke self-tests.

Administrators in any of these roles can login to the TOE either locally or remotely. Because of FIPS level 3 cryptographic ports and interfaces requirements, local administration is permitted only for the purpose of placing the TOE into the evaluated configuration. Use of the serial console for administration is not included in the evaluated configuration, - the console is only allowed for the receipt of alarms, given it is directly connected via a dedicated connection, and given that this connection is physically protected against tampering as is the TOE itself. 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 security administrator is allowed to perform the following:

- 1) specify the interval at which the TSF self tests periodically run
- 2) enable, disable, determine and modify the set of rules that indicates potential violations (FAU_SAA)
- 3) enable, disable, determine and modify the set of audited events (FAU_SEL)
- 4) perform searches and sorting of audit data (FAU_SAR)
- 5) manipulate the security attributes referenced in the TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, TRANSPARENT MODE FIREWALL SFP, ROUTE MODE FIREWALL SFP and UNAUTHENTICATED TOE SERVICES SFP policies
- 6) query, modify, delete, clear all TSF data, except cryptographic security data and audit data
- 7) enable or disable the audible alarm mechanism on alarm messages (FAU_ARP)
- 8) acknowledge alarm messages (FAU_ARP)
- 9) enable, disable the ICMP functions
- 10) determine and modify the administrator-specified network identifier or set of identifiers that are used for the monitoring of resource utilization quotas (FRU_RSA)
- 11) determine and modify the administrator-specified period of time that is used for the monitoring of resource utilization quotas (FRU_RSA)
- 12) specify alternative initial values to override the default values for the TRANSPARENT and ROUTE MODE VPN SFPs
- 13) specify alternative initial values to override the default values for the TRANSPARENT and ROUTE MODE FIREWALL SFPs

14) specify alternative initial values to override the default values for the UNAUTHENTICATED TOE SERVICES SFP and

15) set the time and date used to form the time stamps.

The cryptographic administrator is allowed to perform the following:

- enable or disable the cryptographic and key generation self-tests;
- perform searches and sorting of audit data;
- modify cryptographic security data;
- acknowledge alarm messages (FAU_ARP);
- execute TSF self tests.

The audit administrator is allowed to perform the following:

- perform searches and sorting of audit data;
- acknowledge alarm messages (FAU_ARP);
- delete audit log entries;
- execute TSF self tests.

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

- FMT_MOF.1(1): The security administrator can specify the interval at which the TSF self-tests run.
- FMT_MOF.1(2): The cryptographic administrator can enable or disable the cryptographic and key generation self-tests,
- FMT_MOF.1(3): All administrators are allowed to perform searches and sorting of audit data.
- FMT_MOF.1(4): The security administrator can enable, disable, determine and modify the set of rules that indicates potential violations and the set of audited events.
- FMT_MOF.1(5): The security administrator can enable or disable the audible alarm mechanism on alarm messages.
- FMT_MOF.1(6): The security administrator can enable and disable the ICMP functions.
- FMT_MOF.1(7): The security administrator determines and modifies the administrator-specified network identifier or set of identifiers that are used for the monitoring of resource utilization quotas. The security administrator determines and modifies the administrator-specified period of time that is used for the monitoring of resource utilization quotas. Also, the security administrator configures the quotas for controlled connection-oriented resource allocation (FRU_RSA.1(2)).
- FMT_MSA.1: The security administrator manipulates the security attributes referenced in the TRANSPARENT MODE VPN SFP, ROUTE MODE VPN SFP, TRANSPARENT MODE FIREWALL SFP, ROUTE MODE FIREWALL SFP and UNAUTHENTICATED TOE SERVICES SFP policies.
- FMT_MSA.3 (1): By default, a security appliance denies all traffic in all directions.
- FMT_MSA.3 (2): By default, the UNAUTHENTICATED TOE SERVICES SFP denies all services. If the security administrator enables a service, it will remain enabled across a device reset.
- FMT_MTD.1(1): Only administrators that are given accounts are permitted to login (either locally or remotely) to the TOE.
- FMT_MTD.1(2): The cryptographic administrator is allowed to modify cryptographic data within the TOE.
- FMT_MTD.1(3): Only the security administrator can use the 'set time' command to modify the time used by the TOE. An authorized NTP Server can also affect the time used by the TOE.
- FMT_MTD.1(4): Only the Security Administrator is allowed to query, modify, delete and create the VPN Policy rules. The command line interface functions used by the security administrator are the 'set vpn' and 'get vpn' commands.

- FMT_MTD.2(1): The TSF restricts the specification of the limits for quotas on transport-layer connections, i.e. TCP Synflood protection, to the Security Administrator. The TSF drops traffic exceeding the limits, logging that it does so.
- FMT_MTD.2(2): The TSF restricts the specification of the limits for quotas on controlled connection-oriented resources to the Security Administrator. The SA may configure the TSF to either deny and log the traffic, or allow and log. The default behavior is to deny and log.
- FMT_REV.1: All configuration changes made to the TOE become effective immediately after the command has succeeded. This occurs even though the configuration change itself may not yet have been saved to permanent storage. This includes revocation of an administrative role, changes to the information flow policy ruleset, disabling a service to unauthenticated users and changes to security associations.
- 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.2: 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. As mentioned above, administration via local console is not permitted by FIPS 140-2 level 3 cryptographic ports and interfaces requirements. However, the local console may be used for monitoring alarms.

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_RCV.1:** Security appliances provide the capability to recover the configuration of an appliance to a Last-Known-Good configuration. Therefore in the event that an appliance has been reconfigured to an inoperable or vulnerable state, the appliance can be quickly recovered to its Last-Known-Good configuration. When a failure or service discontinuity occurs the TOE automatically enters maintenance mode (i.e., the crash dump screen). From this mode an administrator can recover to the Last-Known-Good configuration from within the console. This functionality, however, requires that an administrator has established a recovery point as a Last-Known-Good configuration.
- **FPT_RPL.1:** The TSF has the ability to perform a preconfigured action that is selected by the security administrator when a replay attack is detected. The configurable actions include rejection of session establishment in case of an administrative session. By default all the replay events are logged along with the timestamp and source address. The TSF provides replay detection on all the physical interfaces and tunnel interfaces as well. For the remote administration sessions, the SSH daemon has a replay detection mechanism.
- **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.
- **FPT_TST_(EXT).1:** The TSF runs self tests during initial start-up. Also, the security administrator can schedule the self tests to be executed or can invoke the self tests at any point during normal operation of the TOE. The TSF provides a command line operation that only the administrator can use to verify the integrity of the static code (including the TSF executable code). For every command line operation that is entered at the command line interface, there is a command interpreter module which checks for the user's privilege levels in real time before an action is taken. These tests demonstrate the integrity of the TSF executable code.
- **FPT_TST.1(1):** In FIPS mode, the TSF provides a suite of cryptographic module self-tests which are invoked every time the device boots up. These crypto self tests can be executed on demand only by the crypto administrator and as scheduled by the security administrator. Also, the TSF can be configured such that the self tests are run automatically after every key is generated. By default this is turned off and the administrator has to turn this on. These tests demonstrate the correct operation of the key error detection mechanism, the cryptographic algorithms, and the RNG/PRNG mechanism as well as verifying the integrity of TSF data related to cryptography.
- **FPT_TST.1(2):** In FIPS mode, the TSF provides a suite of cryptographic module self-tests which are invoked every time the device boots up. These crypto self tests can be executed on demand only by the crypto administrator and as scheduled by the security administrator. Also, the TSF can be configured such that the self tests are run automatically after every key is generated. By default this is turned off and the administrator has to turn this on. These tests demonstrate the correct operation of the TSF's key generation mechanism by verifying the integrity of the TSF executable code that performs key generation. The tests also verify the integrity of TSF data related to key generation.

6.1.7 Resource utilization

The Resource utilization security function is described in the context of how it satisfies the Resource utilization security requirements.

The Resource utilization function is designed to satisfy the following security functional requirements:

- **FRU_RSA.1(1):** The TSF represents a transport-layer communication pathway as a connection. The TOE can defend itself and the resources it protects from various DoS and DDoS attacks by rate limiting these connections (i.e., applying a maximum quota to the number of connections). TCP SYN flood attack protection is one of them.

- FRU_RSA.1(2): The TSF utilizes a session table to control connection oriented resources. The TSF supports both source based session limiting and destination based session limiting to prevent session table flooding attacks. The thresholds for the same are configurable by the security administrator. Also, the Security Administrator has the ability to define the maximum number of resources a particular address or set of addresses can use over a specified time period. The TSF also provides SYN-ACK-ACK proxy flood, SYN flood and SYN cookie protection mechanisms. The action taken by the TSF when a specific IT entity exceeds the quota is reject and log.

6.1.8 TOE access

Due to restrictions imposed by FIPS certification the use of a local console, except for monitoring alarms, is not permitted when the system is in FIPS mode. However, the TSF provides functionality at the local console that supports the TOE access security function as described in the following bullets. That is, the TOE functionality is described in the context of how it satisfies the TOE access security requirements.

The TOE access function is designed to satisfy the following security functional requirements:

- FTA_SSL.1: The TSF overwrites the display device and makes the current contents unreadable after the local interactive session is locked due to inactivity, thus disabling any further interaction with the TOE. This mechanism is the inactivity timer for administrative sessions. The security administrator can configure this inactivity timer on administrative sessions after which the session will be locked. ScreenOS requires the administrative user to re-authenticate after a timeout to unlock the session.
- FTA_SSL.2: The local administrative user can logout of an existing session by typing 'logout' to exit the CLI admin session and the TSF makes the current contents unreadable after the admin initiates the locking and no user activity can take place until the user re-identifies and authenticates. The TOE does not support background processes.
- FTA_SSL.3: The TSF terminates all remote administrator sessions after a security administrator configured time interval of session inactivity has elapsed.
- FTA_TAB.1: The TSF displays the banner that is set by the security administrator before the prompt for identification and authentication appears on the screen. The contents of the banner are modifiable by the security administrator including the firmware version number and model number.
- FTA_TSE.1: The security administrator can restrict the establishment of an administrative session based on a schedule (i.e., day and time) or based upon the originating source ip address (or subnet). These restrictions are unique to each administrator.

6.1.9 Trusted path/channels

The trusted communication channel between the TSF and an entity in the environment is provided by the use of a certificate based IPSEC VPN tunnel. The use of public key based certificates for establishing the IPSEC tunnel provides assured identification of the end points. The IPSEC VPN tunnel provides the encryption for protection of the channel data from disclosure and integrity for detection of the modification of channel data.

After the establishment of the IPSEC VPN tunnel, either of the end points (i.e. the TSF or the environment) can initiate the communication via the trusted channel (IPSEC VPN tunnel). Authentication attempts over a trusted channel occur only in the context of remote administration, where a login attempt is being made by an administrator. When using the command line interface to perform a login, the administrator identifies themselves and the TSF prompts for authentication data. Thus, the TSF initiates communication via the trusted channels for all authentication attempts.

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, assured identification is guaranteed by using public key certificate based authentication for SSH. The SSH protocol ensures that the data transmitted over a SSH session cannot be disclosed or altered.

The Trusted path/channels function is designed to satisfy the following security functional requirements:

- FTP_ITC.1(1): The trusted communication channel between the TSF and an authorized IT entity in the environment is provided by the use of a certificate based IPSEC VPN tunnel that utilizes the FIPS cryptographic module. The IPSEC VPN tunnel provides the protection for the channel data from disclosure. All trusted channels to authorized servers must be initiated by the TOE. The TOE acts solely as a client to NTP servers and does not act as an NTP server itself.
- FTP_ITC.1(2): The trusted communication channel between the TSF and an authorized IT entity in the environment is provided by the use of a certificate based IPSEC VPN tunnel that utilizes the FIPS cryptographic module. The TSF supports the security services of integrity validation (i.e., detection of the modification of data) as defined by IPSEC. All trusted channels to authorized servers must be initiated by the TOE. .
- FTP_TRP.1(1): 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. ScreenOS ensures that administrators must use SSH to protect all remote administration sessions.
- FTP_TRP.1(2): The encrypted communication channel between the TSF and an administrator in the environment is provided by the use of a SSH session using public key certificate based authentication. This protocol provides integrity of the transmitted data that utilizes the FIPS cryptographic module. ScreenOS ensures that administrators must use SSH to protect all remote administration sessions.

7. Protection Profile Claims

The TOE conforms to the following protection profiles:

- U.S. Government Virtual Private Network (VPN) Boundary Gateway Protection Profile For Medium Robustness Environments, Version 1.2, January 2009.
- U.S. Government Traffic-Filter Firewall Protection Profile For Medium Robustness Environments, Version 1.1, July 25, 2007.

NOTE: The testing associated with AVA_CCA_(EXT).1 and AVA_VAN.4 as required by these Protection Profiles is pending completion by NSA.

This security target uses a slightly different notation for labeling the elements of a requirement than that notation used by the protection profiles. That difference is the location of the element number in the requirement name. Within the protection profiles the element number occasionally can be found within the middle of the requirement number scheme. This usually occurs when a requirement was being interpreted. For example, for the requirement that incorporates the NIAP interpretation of FAU_GEN.1, the VPN PP includes a requirement identified as FAU_GEN.1-NIAP-0410. The VPN PP then denotes the two elements for this requirement as FAU_GEN.1.1-NIAP-0407 and FAU_GEN.1.2-NIAP-0407. The ST Author chose to use a consistent notation of elements within requirements that always placed the element number at the end of the identifier. Thus, the notation used within this ST differs from that of the PPs that it claims conformance to, but the meaning of requirements is unchanged. The following are some typical examples of notation differences. They are not meant to be a complete list.

Element Identifier from the VPN PP	Element Identifier in this ST
FAU_GEN.1.2-NIAP-0407	FAU_GEN.1-NIAP-0410.2
FAU_SEL.1.1-NIAP-0407	FAU_SEL.1-NIAP-0407.1

7.1 Security Environment

All of the security environment and objective statements have been drawn from a validated PP (Medium Robustness VPN Protection Profile or the Medium Robustness TFFW Protection Profile). Please consult those PPs for the applicable correspondence rationale. This ST includes all of the Threats, Policies and Assumptions from both protection profiles.

Table 3 PP Correspondence Rationale for Threats

Threat	Threat in TFFW PP	Threat in VPN PP	Threat in this ST
T.ADDRESS_MASQUERADE	Same as ST	Same as ST	Matches Both PP
T.ADMIN_ERROR	Same as ST	Same as ST	Matches Both PP
T.ADMIN_ROGUE	Same as ST	Same as ST	Matches Both PP
T.AUDIT_COMPROMISE	Same as ST	Same as ST	Matches Both PP
T.CRYPTO_COMPROMISE	Same as ST	Same as ST	Matches Both PP
T.FLAWED_DESIGN	Same as ST	Same as ST	Matches Both PP
T.FLAWED_IMPLEMENTATION	Same as ST	Same as ST	Matches Both PP
T.MALICIOUS_TSF_COMPROMISE	Same as ST	Same as ST	Matches Both PP
T.MASQUERADE	Same as ST	Same as ST	Matches Both PP
T.POOR_TEST	Same as ST	Same as ST	Matches Both PP
T.REPLAY	Same as ST	Same as ST	Matches Both PP
T.RESIDUAL_DATA	Same as ST	Same as ST	Matches Both PP
T.RESOURCE_EXHAUSTION	Same as ST	Same as ST	Matches Both PP
T.SPOOFING	Same as ST	Same as ST	Matches Both PP

T.UNATTENDED_SESSION	Same as ST	Same as ST	Matches Both PP
T.UNAUTHORIZED_ACCESS	Same as ST	Same as ST	Matches Both PP
T.UNAUTHORIZED_PEER	No	Same as ST	Matches VPN PP
T.UNIDENTIFIED_ACTIONS	Same as ST	Same as ST	Matches Both PP
T.UNKNOWN_STATE	Same as ST	Same as ST	Matches Both PP

T.UNAUTHORIZED_PEER An unauthorized IT entity may attempt to establish a security association with the TOE.

Table 4 PP Correspondence Rationale for Security Policies

Policies	Threat in TFFW PP	Threat in VPN PP	Threat in this ST
P.ACCESS_BANNER	Same as ST	Same as ST	Matches Both PP
P.ACCOUNTABILITY	Same as ST	Same as ST	Matches Both PP
P.ADMIN_ACCESS	Same as ST	Same as ST	Matches Both PP
P.CRYPTOGRAPHIC_FUNCTIONS	Same as ST	Same as ST	Matches Both PP
P.CRYPTOGRAPHY_VALIDATED	Same as ST	Same as ST	Matches Both PP
P.INTEGRITY	No	Same as ST	Matches VPN PP
P.VULNERABILITY_ANALYSIS_TEST	Same as ST	Same as ST	Matches Both PP

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).

Table 5 PP Correspondence Rationale for Assumptions

Policies	Policies in TFFW PP	Policies in VPN PP	Policies in this ST
A.NO_GENERAL_PURPOSE	Assumes no general purpose computing capabilities on the TOE.	Assumes that administrator configures system such that no general purpose computing capabilities are available on the TOE.	Matches VPN PP
A.PHYSICAL	Same as ST	Same as ST	Matches Both PP
A.NO_TOE_BYPASS	Same as ST	Same as ST	Matches Both PP

Table 6 PP Correspondence Rationale for Security Objectives

Objectives	Objectives in TFFW PP	Objectives in VPN PP	Objectives in this ST
O.ADMIN_ROLE	Objective in this PP does not include the phrase, "and to make the administrative functions available locally and remotely".	Same as ST	Matches VPN PP
O.AUDIT_GENERATION	Same as ST	Same as ST	Matches Both PP
O.AUDIT_PROTECTION	Same as ST	Same as ST	Matches Both PP
O.AUDIT_REVIEW	Same as ST	Same as ST	Matches Both PP
O.CHANGE_MANAGEMENT	Same as ST	Same as ST	Matches Both PP
O.CORRECT_TSF_OPERATION	Same as ST	Same as ST	Matches Both PP
O.CRYPTOGRAPHIC_FUNCTIONS	Same as ST	Same as ST	Matches Both PP
O.CRYPTOGRAPHY_VALIDATED	Same as ST	Same as ST	Matches Both PP
O.DISPLAY_BANNER	Same as ST	Same as ST	Matches Both PP
O.DOCUMENT_KEY_LEAKAGE	Same as ST	Same as ST	Matches Both PP
O.INTEGRITY	No Matching Objective	Same as ST	Matches VPN PP
O.MAINT_MODE	Same as ST	Same as ST	Matches Both PP
O.MANAGE	Same as ST	Same as ST	Matches Both PP
O.MEDIATE	Same as ST	Same as ST	Matches Both PP

O.PEER_AUTHENTICATION	No Matching Objective	Same as ST	Matches VPN PP
O.REPLAY_DETECTION	Same as ST	Same as ST	Matches Both PP
O.RESIDUAL_INFORMATION	Same as ST	Same as ST	Matches Both PP
O.RESOURCE_SHARING	Similar but uses different examples	Same as ST	Matches VPN PP
O.ROBUST_ADMIN_GUIDANCE	No Matching Objective	Same as ST	Matches VPN PP
O.ROBUST_TOE_ACCESS	Same as ST	Same as ST	Matches Both PP
O.SELF_PROTECTION	Same as ST	Same as ST	Matches Both PP
O.SOUND_DESIGN	Same as ST	Same as ST	Matches Both PP
O.SOUND_IMPLEMENTATION	Same as ST	Same as ST	Matches Both PP
O.THOROUGH_FUNCTIONAL_TESTING	Same as ST	Same as ST	Matches Both PP
O.TIME_STAMPS	Same as ST	Same as ST	Matches Both PP
O.TRUSTED_PATH	Trusted Path needed for admin	Trusted path needed for all users	Matches TFFW PP
O.VULNERABILITY_ANALYSIS	Same as ST	Same as ST	Matches Both PP

Table 7 PP Correspondence Rationale for Objectives for the Environment

Objectives for the Environment	Objectives in TFFW PP	Objectives in VPN PP	Objectives in this ST
OE.CRYPTANALYTIC	Same as ST	Same as ST	Matches Both PP
OE.NO_GENERAL_PURPOSE	Same as ST	Same as ST	Matches Both PP
OE.NO_TOE_BYPASS	Same as ST	Same as ST	Matches Both PP
OE.PHYSICAL	Same as ST	Same as ST	Matches Both PP

7.2 Security Functional Requirements

The following table summarizes how the requirements from this security target correspond to the requirements from the Medium Robustness VPN Protection Profile and the Medium Robustness TFFW Protection Profile.

Table 8 PP Correspondence Rationale for SFRs

Security Target SFR	Traffic-Filter FW PP	VPN PP
FAU_ARP.1	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FAU_ARP_ACK_(EXT).1	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FAU_GEN.1-NIAP-0410	The events from this PP are all reproduced in the ST.	The VPN PP contained a different interpretation (FAU_GEN.1-NIAP-0407) than the one used in the TFFW PP. The text of the requirement as completed by this PP requires more detail within each audit record (i.e., this PP requires a subject identity when applicable) than the text from the requirement FAU_GEN.1-NIAP-0407 from this the VPN PP. Therefore, the interpretation used by the TFFW PP is being included in this ST. Furthermore, the events from this PP are all reproduced in the ST.
FAU_GEN.2-NIAP-0410	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FAU_SAA.1-NIAP-0407	ST includes all monitoring rules found in the PP.	ST includes all monitoring rules found in the PP.
FAU_SAR.1	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FAU_SAR.2	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FAU_SAR.3	ST includes all searching and sorting criteria	ST includes all searching and sorting criteria

	found in the PP.	found in the PP.
FAU_SEL.1-NIAP-0407	ST includes all selectable attributes found in the PP.	ST includes all selectable attributes found in the PP.
FAU_STG.1-NIAP-0429	The PP attempted to modify the FAU_STG.1 requirement from Version 3.1 of the CC. However, the removal of “unauthorised” from the second element, contradicts the exception explicitly permitted by the first element. Therefore, the FAU_STG.1 requirement in the TFFW PP is not sound. Since the first element is explicitly changed to describe an ‘authorized’ deletion, the version of this SFR from the VPN PP is used because it is more internally consistent and has the same meaning.	Requirement in PP is the same as this ST
FAU_STG.3	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FAU_STG.NIAP-0414-1-NIAP-0429	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_BCM_(EXT).1	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_CKM.1(1) Cryptographic key generation (for symmetric keys)	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_CKM.1(2) Cryptographic key generation (for asymmetric keys)	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_CKM.2 Cryptographic Key Distribution	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_CKM_(EXT).2 Crypto Key Handling and Storage	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_CKM.4	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_COP.1(1)	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_COP.1(2)	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_COP.1(3)	This ST includes a version of this requirement as stated in the VPN PP. The requirements use different wording but both require SHA-256.	Requirement in PP is the same as this ST
FCS_COP.1(4)	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_COP_(EXT).1	Requirement in PP is the same as this ST	Requirement in PP is the same as this ST
FCS_IKE_(EXT).1	No matching requirement	Requirement in PP is the same as this ST
FDP_IFC.1(1) & FDP_IFC.1(2)	No matching requirement	This ST includes 2 iterations of FDP_IFC.1(1) from the VPN PP to describe the transparent and route modes supported by this product.
FDP_IFC.1(3)	ST renames FDP_IFC.1(2)	ST renames FDP_IFC.1(2)
FDP_IFC.1(4) & FDP_IFC.1(5)	This ST includes 2 iterations of FDP_IFC.1(1) from the TFFW PP to describe the transparent and route modes supported by this product.	No matching requirement
FDP_IFF.1.1(1) & FDP_IFF.1.1(2)	No matching requirement	This ST includes 2 iterations of FDP_IFF.1.1(1) from the VPN PP to describe the transparent and route modes supported by

		this product.
FDP_IFF.1.1(3)	ST renames FDP_IFF.1.1(2) from the TFFW PP	ST renames FDP_IFF.1.1(2) from the VPN PP
FDP_IFF.1.1(4) & FDP_IFF.1.1(5)	This ST includes 2 iterations of FDP_IFF.1.1(1) from the TFFW PP to describe the transparent and route modes supported by this product.	No matching requirement
FDP_RIP.2	Same as ST	Same as ST
FIA_AFL.1	This SFR from TFFW PP was modified to indicate that the integer must be positive. The VPN PP included this limitation and since having a configuration value here with a negative value was infeasible, the limitation from the VPN PP was included in this ST.	The SFR from the VPN PP required the failure handling be “administrator” configurable, while the TFFW PP indicated this be configurable by the “security administrator”. This ST used the more restrictive text.
FIA_ATD.1	The TFFW PP used the term “authorized user”, while the VPN PP used the term “administrator”. However the application note in both PP’s indicated that the requirement was intended to be scoped to the same set of users. Since no difference in meaning was intended by the PP authors, either verbiage could be used and the verbiage from the VPN PP was chosen.	Same as ST
FIA_UAU.1 (1)	Same as ST	Same as ST
FIA_UAU.2	The ST requirement matches FIA_UAU_EXP.2 from the TFFW PP.	The VPN PP requires authentication for “the administrator” prior to TSF-mediated actions. The TFFW PP requires authentication for “administrators and authorized IT entities” prior to TSF-mediated actions. Since the requirement in the TFFW PP is requiring more authentication actions, that version of the requirement is being included in this ST.
FIA_UAU_(EXT).5	Same as ST	Same as ST
FIA_UID.2	Same as ST	Same as ST
FIA_USB.1	Same as ST	Same as ST
FMT_MOF.1(1)	The SFR contained a reference to a specific self testing function as defined by another SFR. The reference was updated by this ST to refer to the function by the SFR identifier excluding its text name (which was incorrect in the PPs). This approach was intended to more accurately identify the self testing function being required.	The SFR contained a reference to a specific self testing function as defined by another SFR. The reference was updated by this ST to refer to the function by the SFR identifier excluding its text name (which was incorrect in the PPs). This approach was intended to more accurately identify the self testing function being required.
FMT_MOF.1(2)	Otherwise, the requirement in this ST is intended to be the same as the requirement in the PP. The SFR contained a reference to a specific self testing function as defined by another SFR. The reference was updated by this ST to refer to the function by the SFR identifier excluding its text name (which was incorrect in the PPs). This approach was intended to more accurately identify the self testing function being required.	Otherwise, the requirement in this ST is intended to be the same as the requirement in the PP. The SFR contained a reference to a specific self testing function as defined by another SFR. The reference was updated by this ST to refer to the function by the SFR identifier excluding its text name (which was incorrect in the PPs). This approach was intended to more accurately identify the self testing function being required.
	Otherwise, the requirement in this ST is	Otherwise, the requirement in this ST is

	intended to be the same as the requirement in the PP.	intended to be the same as the requirement in the PP.
FMT_MOF.1(3), FMT_MOF.1(4), FMT_MOF.1(5), & FMT_MOF.1(6) FMT_MOF.1(7)	Same as ST	Same as ST
FMT_MSA.1	Same as ST	All of the functions mentioned in the iteration of this SFR in the VPN PP are included in the iteration of this requirement used in this ST. The VPN PP allows “an Administrator” to specify security policy attributes, while the TFFW PP allowed only the “Security Administrator”. The ST choose the most restrictive version allowing only the “Security administrator.”
FMT_MSA.3 (1)	The version of this SFR from the TFFW PP (i.e., FMT_MSA.3-NIAP-0409(1)) removed some of the words from the original CC requirement the meaning appears to be consistent with the SFR from the VPN PP. Therefore, the version from the VPN PP (the on more closely reflecting the original CC SFR is used. The requirement in this ST also references all of the policies from this ST that correspond to the policies referred to by this SFR in the VPN and the TFFW PPs.	The requirement in this ST references all of the policies from this ST that correspond to the policies referred to by this SFR in the VPN and the TFFW PPs. Other than the addition of policies from the TFFW PP, this requirement is the same in this ST as it is in the VPN PP.
FMT_MSA.3(2)	The version of this SFR from the TFFW PP (i.e., FMT_MSA.3-NIAP-0409(2)) removed some of the words from the original CC requirement the meaning appears to be consistent with the SFR from the VPN PP. Therefore, the version from the VPN PP (the on more closely reflecting the original CC SFR is used.	Same as ST
FMT_MTD.1(1)	Same as ST	Same as ST
FMT_MTD.1(2)	Same as ST	Same as ST
FMT_MTD.1(3)	Same as ST	Same as ST
FMT_MTD.1(4)	This ST changes the list of applicable rule sets to include a list all SFP rules that are included in this ST.	This ST changes the list of applicable rule sets to include a list all SFP rules that are included in this ST.
FMT_MTD.2(1)	Same as ST	Same as ST
FMT_MTD.2(2)	Same as ST	Same as ST
FMT_REV.1	Same as ST	Same as ST
FMT_SMF.1	No matching requirement	Same as ST
FMT_SMR.2	Same as ST	Same as ST
FPT_RCV.1	Same as ST	Same as ST
FPT_RPL.1	Same as ST	Same as ST
FPT_STM.1	Same as ST	Same as ST
FPT_TST_(EXT).1	Same as ST	Same as ST
FPT_TST.1(1)	This requirement was updated in the ST to	This requirement was updated in the ST to

	reference material within the ST rather than material in the protection profile. The SFR referred to “Appendix A”, while this ST refers to “Section 10”. “Appendix A” from the PP was duplicated in “Section 10” of this ST. The change simply clarifies the reference.	reference material within the ST rather than material in the protection profile. The SFR referred to “Appendix F”, while this ST refers to “Section 10”. “Appendix F” from the PP was duplicated in “Section 10” of this ST. The change simply clarifies the reference.
FPT_TST.1(2)	Same as ST	Same as ST
FRU_RSA.1(1)	Same as ST	The VPN PP enforces quotas upon users, which are really just known “source subject identifiers”. Therefore, the terminology from the TFFW PP has been used.
FRU_RSA.1b	Same as ST	Same as ST
FTA_SSL.1	This SFR uses different wording in the TFFW and VPN PP, however, both requirements have the same meaning. Therefore, the version of this requirement that was used in this security target matches that used in the VPN PP because it more closely matches the generic CC requirement.	Same as ST
FTA_SSL.2	This SFR uses different wording in the TFFW and VPN PP, however, both requirements have the same meaning. Therefore, the version of this requirement that was used in this security target matches that used in the VPN PP because it more closely matches the generic CC requirement.	Same as ST
FTA_SSL.3	Same as ST	Same as ST
FTA_TAB.1	Same as ST	Same as ST
FTA_TSE.1	The TFFW PP refined this SFR to “an authorized user session, while the VPN PP refers to “an administrator session”. Since administrators are the only authorized users in the TOE, this ST choose to use the wording from the VPN PP.	Same as ST
FTP_ITC.1(1)	Same as ST	Same as ST
FTP_ITC.1(2)	Same as ST	Same as ST
FTP_TRP.1(1)	The TFFW PP refers to “remote users” while the VPN PP refers to “administrators” in FTP_TRP.1.2(1). Since these refer to the same set of accounts, the wording from the VPN PP was chosen for use in this ST.	Same as ST
FTP_TRP.1(2)	The TFFW PP refers to “remote users” while the VPN PP refers to “administrators” in FTP_TRP.1.2(1). Since these refer to the same set of accounts, the wording from the VPN PP was chosen for use in this ST.	Same as ST

7.3 Assurance Requirements

The assurance requirements in the VPN PP and the TFFW PP are identical.

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 have been drawn from validated PPs (Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile). Please consult those PPs for the applicable correspondence rationale.

8.2 Security Requirements Rationale

All of the security functional requirements have been drawn from validated PPs (Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile). Please consult those PPs for the applicable rationale.

8.3 Requirement Dependency Rationale

All of the requirements have been drawn from a validated PP (Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile). Please consult those PPs for the applicable rationale.

8.4 Extended Requirements Rationale

There are extended requirements in this Security Target. All of the extended requirements levied upon the TOE have been drawn from a validated PP (Medium Robustness VPN Protection Profile and Medium Robustness TFFW Protection Profile). Please consult those PPs for the rationale for these extended requirements.

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 9 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	Resource utilization	TOE access	Trusted path/channels
FAU_ARP.1	X								
FAU_ARP_ACK_(EXT).1	X								
FAU_GEN.1-NIAP-0410	X								
FAU_GEN.2-NIAP-0410	X								
FAU_SAA.1-NIAP-0407	X								
FAU_SAR.1	X								
FAU_SAR.2	X								
FAU_SAR.3	X								
FAU_SEL.1-NIAP-0407	X								
FAU_STG.1-NIAP-0429	X								
FAU_STG.3	X								
FAU_STG.NIAP-0414-1-NIAP-0429	X								
FCS_BCM_(EXT).1		X							
FCS_CKM.1(1)		X							
FCS_CKM.1(2)		X							
FCS_CKM.2		X							
FCS_CKM_(EXT).2		X							
FCS_CKM.4		X							
FCS_COP.1 (1)		X							
FCS_COP.1(2)		X							
FCS_COP.1(3)		X							
FCS_COP.1(4)		X							
FCS_COP_(EXT).1		X							
FCS_IKE_(EXT).1		X							
FDP_IFC.1(1)			X						
FDP_IFC.1(2)			X						
FDP_IFC.1(3)			X						
FDP_IFC.1(4)			X						
FDP_IFC.1(5)			X						
FDP_IFF.1(1)			X						
FDP_IFF.1(2)			X						
FDP_IFF.1(3)			X						
FDP_IFF.1(4)			X						
FDP_IFF.1(5)			X						
FDP_RIP.2			X						
FIA_AFL.1				X					
FIA_ATD.1				X					
FIA_UAU.1				X					
FIA_UAU.2				X					
FIA_UAU_(EXT).5				X					
FIA_UID.2				X					
FIA_USB.1				X					
FMT_MOF.1(1)					X				
FMT_MOF.1(2)					X				
FMT_MOF.1(3)					X				
FMT_MOF.1(4)					X				
FMT_MOF.1(5)					X				

FMT_MOF.1(6)					X				
FMT_MOF.1(7)					X				
FMT_MSA.1					X				
FMT_MSA.3(1)					X				
FMT_MSA.3(2)					X				
FMT_MTD.1(1)					X				
FMT_MTD.1(2)					X				
FMT_MTD.1(3)					X				
FMT_MTD.1(4)					X				
FMT_MTD.2(1)					X				
FMT_MTD.2(2)					X				
FMT_REV.1					X				
FMT_SMF.1					X				
FMT_SMR.2					X				
FPT_RCV.1						X			
FPT_RPL.1						X			
FPT_STM.1						X			
FPT_TST_(EXT).1						X			
FPT_TST.1(1)						X			
FPT_TST.1(2)									
FRU_RSA.1(1)							X		
FRU_RSA.1(2)							X		
FTA_SSL.1								X	
FTA_SSL.2								X	
FTA_SSL.3								X	
FTA_TAB.1								X	
FTA_TSE.1								X	
FTP_ITC.1(1)									X
FTP_ITC.1(2)									X
FTP_TRP.1(1)									X
FTP_TRP.1(2)									X

Table 9 Security Functions vs. Requirements Mapping

8.6 PP Claims Rationale

See Section 7, Protection Profile Claims.

9. Audit Events

Requirement	Auditable Events	Audit Record Contents
FAU_ARP.1	Potential security violation was detected	Identification of what caused the generation of the alarm
FAU_ARP_ACK_(EXT).1	Acknowledgement of the alarm by an administrator	The identity of the administrators that acknowledged the alarm.
FAU_GEN.1-NIAP-0410	None	
FAU_GEN.2-NIAP-0410	None	
FAU_SAA.1-NIAP-0407	Enabling and disabling of any of the analysis mechanisms	The identity of the Security Administrator performing the function
FAU_SAR.1	Opening the audit trail	The identity of the Audit Administrator performing the function
FAU_SAR.2	Unsuccessful attempts to read information from the audit records	The identity of the administrator performing the function
FAU_SAR.3	None	
FAU_SEL.1-NIAP-0407	All modifications to the audit configuration that occur while the audit collection functions are operating	The identity of the Security Administrator performing the function
FAU_STG.1-NIAP-0429	None	
FAU_STG.3	Actions taken due to exceeding the audit threshold	The identity of the Security Administrator performing the function
FAU_STG.NIAP-0414-1-NIAP-0429	Actions taken due to the audit storage failure	The identity of the Security Administrator performing the function
FCS_BCM_(EXT).1	None	
FCS_CKM.1(1)	Generation and loading of key. Failure of the activity.	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information

FCS_CKM.1(2)	Generation and loading of key. Failure of the activity.	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_CKM.2	Failure of the key distribution	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_CKM.4	Failure of the key distribution	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_CKM.(EXT).2	Failure of the cryptographic key validation and packing	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_COP_(EXT).1	Failure of cryptographic operation	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_COP.1(1)	Failure of cryptographic operation (for data encryption/decryption)	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_COP.1(2)	Failure of cryptographic operation (for cryptographic signature)	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_COP.1(3)	Failure of cryptographic operation (cryptographic hashing)	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_COP.1(4)	Failure of cryptographic operation (for cryptographic key agreement)	Type of cryptographic operation Any applicable cryptographic mode(s) of operation, excluding any sensitive information
FCS_IKE_(EXT).1	<ul style="list-style-type: none"> • Generation and loading of key pair for digital signatures. • Changes to the pre-shared key used for authentication • All modifications to the key lifetimes. • Failure of the authentication in 	If failure occurs, record an English description for the failure.

	<p>Phase 1.</p> <ul style="list-style-type: none"> • Failure to negotiate a security association in Phase 2. 	
FDP_IFC.1 (all iterations)	None	
FDP_IFF.1(1), FDP_IFF.1(2)	<ul style="list-style-type: none"> • Decisions to permit or deny information flows. • Operation applied to each information flow permitted. 	<p>Presumed identity of source subject</p> <p>Identity of destination subject</p> <p>Transport layer protocol, if applicable</p> <p>Source subject service identifier, if applicable</p> <p>Destination subject service identifier, if applicable</p> <p>Identity of the interface on which the TOE received the packet</p> <p>Identity of the rule that allowed or disallowed the packet flow</p> <p>For denied information flows, the reason for denial</p>
FDP_IFF.1(3)	Decisions to permit/deny information flows between a subject and the TOE	<p>Presumed identity of source subject</p> <p>Identity of destination subject</p> <p>Transport layer protocol, if applicable</p> <p>Source subject service identifier, if applicable</p> <p>Destination subject service identifier, if applicable</p> <p>Identity of the interface on which the TOE received the packet</p> <p>Identity of the rule that allowed or disallowed the packet flow</p> <p>For denied information flows, the reason for denial</p>
FDP_IFF.1(4), FDP_IFF.1(5)	<p>Decisions to permit/deny information flows</p> <p>Failure to reassemble fragmented packets</p>	<p>Presumed identity of source subject</p> <p>Identity of destination subject</p> <p>Transport layer protocol, if applicable</p> <p>Source subject service identifier, if applicable</p> <p>Destination subject service identifier,</p>

		<p>if applicable</p> <p>Identity of the firewall interface associated on which the TOE received the packet</p> <p>Identity of the rule that allowed or disallowed the packet flow</p> <p>Reason why fragmented packets could not be reassembled (i.e., invalid fragment identifier, invalid offset, invalid fragment data length)</p>
FDP_RIP.2	None	
FIA_AFL.1	<p>The reaching of the threshold for the unsuccessful authentication attempts</p> <p>The actions (e.g. disabling of an account) taken</p> <p>The subsequent, if appropriate, restoration to the normal state (e.g. re-enabling of an account)</p>	<p>Identity of the unsuccessfully authenticated user</p> <p>Identity of the unsuccessfully authenticated user and the identity of the administrator performing the function.</p>
FIA_ATD.1	None	
FIA_UAU.1	None	
FIA_UAU.2	Successful and unsuccessful use of authentication mechanisms	Claimed identity of the user using the authentication mechanism
FIA_UAU_(EXT).5	All use of the local authentication mechanism	Claimed identity of the user attempting to authenticate
FIA_UID.2	All use of the user identification mechanism used for authorized users (that is, those that authenticate to the TOE)	Claimed identity of the user using the identification mechanism
FIA_USB.1	Success and failure of binding of user security attributes to a subject	The identity of the user whose attributes are attempting to be bound
FMT_MOF.1(1)	All modifications in the behavior of the functions in the TSF	The identity of the administrator performing the function
FMT_MOF.1(2)	Enabling or disabling of the key-generation self-tests	The identity of the administrator performing the function

FMT_MOF.1(3)	All modifications in the behavior of the functions in the TSF	The identity of the administrator performing the function
FMT_MOF.1(4)	All modifications in the behavior of the functions in the TSF	The identity of the administrator performing the function
FMT_MOF.1(5)	All modifications in the behavior of the functions in the TSF	The identity of the administrator performing the function
FMT_MOF.1(6)	All modifications in the behavior of the functions in the TSF	The identity of the administrator performing the function
FMT_MOF.1(7)	All modifications in the behavior of the functions in the TSF	The identity of the administrator performing the function
FMT_MSA.1	All manipulation of the security attributes	The identity of the administrator performing the function
FMT_MSA.3 (1)	None	
FMT_MSA.3 (2)	None	
FMT_MTD.1(1)	All modifications of the values of TSF data by the administrator	The identity of the administrator performing the function
FMT_MTD.1(2)	All modifications of the values of cryptographic security data by the cryptographic administrator	The identity of the administrator performing the function
FMT_MTD.1(3)	All modifications to the time and date used to form the time stamps by the administrator	The identity of the administrator performing the function
FMT_MTD.1(4)	All modifications to the information flow policy ruleset by the Security Administrator	The identity of the security administrator performing the function
FMT_MTD.2(1)	All modifications of the limits Actions taken when the quota is exceed (include the fact that the quota was exceeded)	The identity of the administrator performing the function
FMT_MTD.2(2)	All modifications of the limits Actions taken when the quota is exceed (include the fact that the quota was exceeded)	The identity of the administrator performing the function

FMT_REV.1	All attempts to revoke security attributes	List of security attributes that were attempted to be revoked The identity of the administrator performing the function
FMT_SMR.2	Modifications to the group of users that are part of a role	User IDs that are associated with the modifications The identity of the administrator performing the function
FPT_RCV.1	The fact that a failure or service discontinuity occurred Resumption of the regular operation	Type of failure or service discontinuity
FPT_RPL.1 (including replay of authentication data notification from the authentication server)	Notification that a replay event occurred	Identity of the user that was the subject of the reply attack
FPT_STM.1	Changes to the time	The identity of the administrator if the change was performed by an administrator or the network identifier of the NTP server if the change was performed from an NTP server.
FPT_TST_(EXT).1	Execution of this set of TSF self tests	The identity of the administrator performing the test, if initiated by an administrator
FPT_TST.1 (1)	Execution of this set of crypto TSF self tests	The identity of the administrator performing the test, if initiated by an administrator
FPT_TST.1 (2)	Execution of this set of key generation self tests	The identity of the administrator performing the test, if initiated by an administrator
FRU_RSA.1(1)	None	
FRU_RSA.1(2)	None	
FTA_SSL.1	a) Locking of an interactive session by the session locking mechanism. b) Successful unlocking of an interactive session. c) Any attempts at unlocking an interactive session.	The identity of the user associated with the session being locked or unlocked
FTA_SSL.2	a) Locking of an interactive session by the session locking mechanism. b) Successful unlocking of an interactive session. c) Any attempts at unlocking an interactive session.	The identity of the user associated with the session being locked or unlocked

FTA_SSL.3	The termination of a remote session by the session locking mechanism	The identity of the user associated with the session that was terminated
FTA_TAB.1	None	
FTA_TSE.1	a) Denial of a session establishment due to the session establishment mechanism. b) All attempts at establishment of a user session.	The identity of the user attempting to establish the session For unsuccessful attempts, the reason for denial of the establishment attempt
FTP_ITC.1(1)	a) All attempted uses of the trusted channel functions. b) Identifier of the initiator and target of all trusted channel functions.	Identification of the initiator and target of all trusted channels
FTP_ITC.1(2)	a) All attempted uses of the trusted channel functions. b) Identifier of the initiator and target of all trusted channel functions.	Identification of the initiator and target of all trusted channels
FTP_TRP.1(1)	a) All attempted uses of the trusted path functions. b) Identification of the user associated with all trusted path invocations, if available.	Identification of the claimed user identity
FTP_TRP.1(2)	a) All attempted uses of the trusted path functions. b) Identification of the user associated with all trusted path invocations, if available.	Identification of the claimed user identity

10. Statistical Random Number Generator Tests

A cryptographic module employing random number generators (RNGs) shall perform the following statistical tests for randomness. A single bit stream of 20,000 consecutive bits of output from each RNG shall be subjected to the following four tests: monobit test, poker test, runs test, and long runs test. (These four tests are simply those that formerly existed as the statistical RNG tests in Federal Information Processing Standard 140-2. However, for purposes of meeting this protection profile, these tests must be performed at the frequency specified earlier in this protection profile.)

The Monobit Test:

1. Count the number of ones in the 20,000 bit stream. Denote this quantity by X.
2. The test is passed if $9,725 < X < 10,275$.

The Poker Test:

1. Divide the 20,000 bit stream into 5,000 contiguous 4 bit segments. Count and store the number of occurrences of the 16 possible 4 bit values. Denote $f(i)$ as the number of each 4 bit value i , where $0 \leq i \leq 15$.
2. Evaluate the following:

$$X = (16/5000) * \left(\sum_{i=0}^{15} [f(i)]^2 \right) - 5000$$

3. The test is passed if $2.16 < X < 46.17$.

The Runs Test:

1. A run is defined as a maximal sequence of consecutive bits of either all ones or all zeros that is part of the 20,000 bit sample stream. The incidences of runs (for both consecutive zeros and consecutive ones) of all lengths (> 1) in the sample stream should be counted and stored.
2. The test is passed if the runs that occur (of lengths 1 through 6) are each within the corresponding interval specified in the table below. This must hold for both the zeros and ones (i.e., all 12 counts must lie in the specified interval). For the purposes of this test, runs of greater than 6 are considered to be of length 6.

Table C.1 - Required Intervals for Length of Runs Test

Length of Run	Required Interval
1	2343 - 2657
2	1135 - 1365
3	542 - 708
4	251 - 373
5	111 - 201
6 and greater	111 - 201

The Long Runs Test:

1. A long run is defined to be a run of length 26 or more (of either zeros or ones).
2. On the sample of 20,000 bits, the test is passed if there are no long runs.