Firewall

FortiOS™ Handbook v3

for FortiOS 4.0 MR3

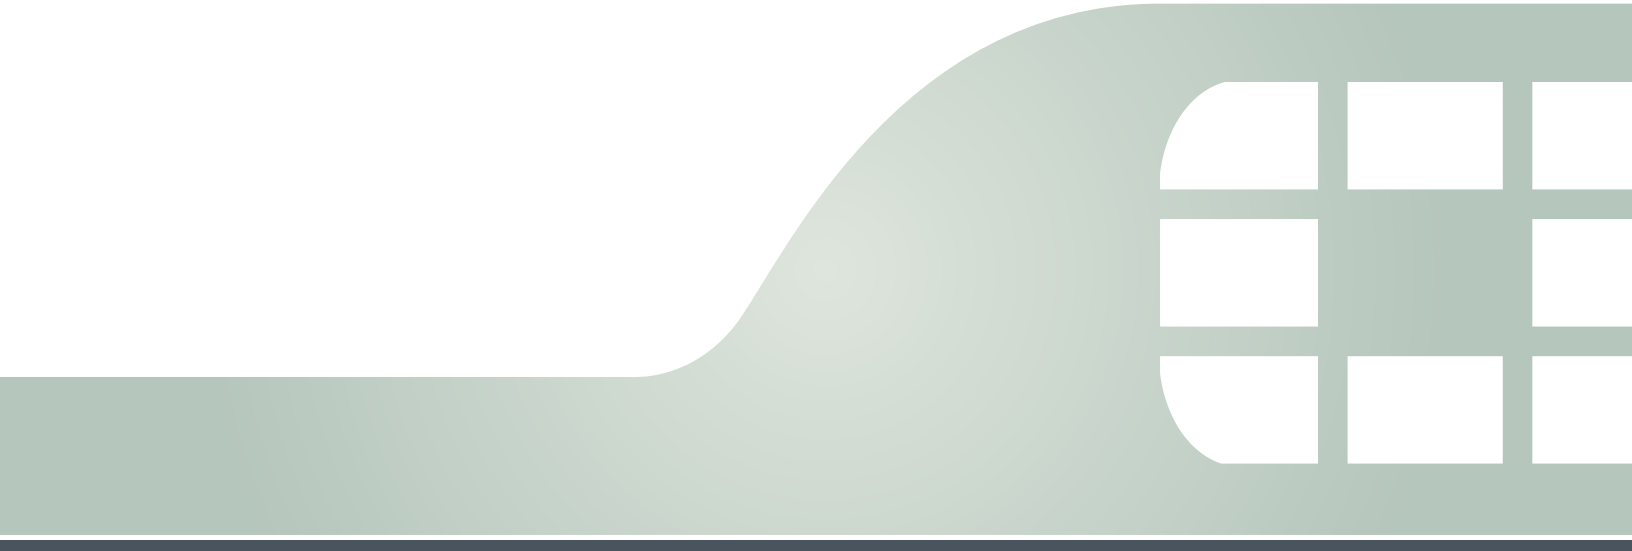

FortiOS™ Handbook Firewall

v3

24 January 2012

01-432-148222-20120124

© Copyright 2012 Fortinet, Inc. All rights reserved. Contents and terms are subject to change by Fortinet without prior notice. Reproduction or transmission of this publication is encouraged.

Trademarks

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Visit these links for more information and documentation for your Fortinet products:

Fortinet Knowledge Base - http://kb.fortinet.com

Technical Documentation - http://docs.fortinet.com

Training Services - http://campus.training.fortinet.com

Technical Support - http://support.fortinet.com

You can report errors or omissions in this or any Fortinet technical document to techdoc@fortinet.com.

Contents

Introduction 9

Before you begin . . . 9
How this guide is organized . . . 9

Understanding the FortiGate firewall 11

What is the FortiGate firewall? . . . 11
FortiGate firewall components . . . 11
How the firewall components create a FortiGate firewall and help in protecting your network . . . 12
Understanding how a packet travels through the FortiGate unit . . . 13
How packets flow in and out of the FortiGate unit . . . 14

Working with NAT in FortiOS 17

NAT in FortiOS . . . 17
NAT/Route mode . . . 17
Route mode . . . 18
Transparent mode . . . 19
Types of NAT in FortiOS . . . 19
Static NAT (SNAT) . . . 20
Static Destination NAT (SDNAT) . . . 20
Static NAT port forwarding . . . 20
Dynamic NAT (DNAT) . . . 20
Dynamic source address translation . . . 21
Dynamic destination address . . . 21
Dynamic port forwarding . . . 21
Combining types of NAT . . . 21

Firewall components 23

Using Interfaces and zones in the FortiGate firewall . . . 23
How to apply VLANs and zones and to a security policy . . . 23
Understanding the firewall address component . . . 24
IP addresses for self-originated traffic . . . 25
IP pools . . . 26
IP Pools for security policies that use fixed ports . . . 27
Source IP address and IP pool address matching . . . 27
Geography-based addressing . . . 28
Wildcard addresses . . . 29

Using wildcard addresses in the firewall configuration . . . 31
Fully Qualified Domain Name addresses . . . 31
Address groups . . . 32
Virtual IP addresses . . . 32
Grouping virtual IPs . . . 32
Match-vip . . . 32
Services . . . 33
Predefined service list . . . 33
Custom service groups . . . 38
Firewall schedules . . . 38
Schedule groups . . . 39
Schedule expiry . . . 39
UTM profiles . . . 39
How to use UTM profiles to monitor and protect your network . . . 40

Security policies 43

Security policy overview . . . 43
Security policy list details . . . 44
Viewing security policies . . . 45
Policy order . . . 45
How to arrange policies . . . 47
Security policies . . . 47
Identity-based policies . . . 47
Identity-based policy example . . . 48
SSL VPN policies . . . 49
IPsec policies . . . 49
Accept policies . . . 50
Deny policies . . . 50
How to allow DNS queries to only one DNS server . . . 50
IPv6 policies . . . 51
Security policy 0 . . . 51
Local-in policies . . . 51
Creating basic security policies . . . 52
How to create a basic security policy for Internet access . . . 52
How to test the basic security policy . . . 53
How to verify if traffic is hitting the basic security policy . . . 53

Monitoring firewall traffic 55

Session tables . . . 55
Viewing session tables in the web-based manager . . . 55
Sessions Monitor . . . 55
Viewing session tables in the CLI . . . 56
Proto_state fields: TCP . . . 57
Proto_state fields: SCTP . . . 58

Proto_state fields: UDP . . . 58
Proto_state field for ICMP . . . 58
Monitoring security policy traffic activity . . . 58

Internet Protocol version 6 (IPv6) 61

What is IPv6? . . . 61
IPv6 in FortiOS . . . 61
Dual stack routing configuration . . . 62
IPv4 tunneling configuration . . . 62
Remotely connecting to an IPv6 network over the Internet . . . 62

Advanced FortiGate firewall concepts 65

Central NAT table . . . 65
Central NAT Table configuration settings . . . 66
Stateful inspection of SCTP traffic . . . 66
Configuring FortiGate SCTP filtering . . . 67
Adding an SCTP custom service . . . 68
Adding an SCTP policy route . . . 68
Changing the session time to live for SCTP traffic . . . 69
Port pairing . . . 69
Blocking port 25 to email server traffic . . . 70
Dedicated traffic . . . 71
Restricting traffic on port 25 . . . 72
Blocking HTTP access by IP . . . 73
ICMP packet processing . . . 74
Adding NAT security policies in Transparent mode . . . 74
Adding a static NAT virtual IP for a single IP address and port . . . 77
Double NAT: combining IP pool with virtual IP . . . 79
Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping . . . 81
Traffic shaping and per-IP traffic shaping . . . 83
Endpoint Security . . . 84
Logging traffic . . . 84
Quality of Service (QoS) . . . 85
Identity-based security policies . . . 85
Identity-based policy positioning . . . 85
Identity-based sub-policies . . . 86

Appendix 87

Document conventions . . . 87
IP addresses . . . 87
Example Network configuration . . . 89
Cautions, Notes and Tips . . . 90
Typographical conventions . . . 90
CLI command syntax conventions . . . 91

Entering FortiOS configuration data . . . 93
Entering text strings (names) . . . 93
Entering numeric values . . . 94
Selecting options from a list . . . 94
Enabling or disabling options . . . 94

Registering your Fortinet product . . . 94

Fortinet products End User License Agreement . . . 95

Training . . . 95

Documentation . . . 95
Fortinet Tools and Documentation CD . . . 95
Fortinet Knowledge Base . . . 95
Comments on Fortinet technical documentation . . . 95

Customer service and technical support . . . 95

Index 97

Introduction

Welcome and thank you for selecting Fortinet products for your network protection. This document describes how to configure the FortiGate firewall on your FortiGate unit. This document also provides advanced firewall concepts.

This chapter contains the following topics:

- Before you begin
- How this guide is organized

Before you begin

Before you begin using this guide, please ensure that:

- You have administrative access to the web-based manager and/or CLI.
- The FortiGate unit is integrated into your network.
- The operation mode has been configured.
- The system time, DNS settings, administrator password, and network interfaces have been configured.
- Firmware, FortiGuard Antivirus and FortiGuard Antispam updates are completed.
- FortiGuard Analysis & Management Service is properly configured.

While using the instructions in this guide, note that administrators are assumed to be super_admin administrators unless otherwise specified. Some restrictions will apply to other administrators.

How this guide is organized

This FortiOS Handbook chapter contains the following sections:

Understanding the FortiGate firewall provides general information about what the FortiGate firewall does, what it is comprised of, and explains how a packet travels through the FortiGate unit.

Working with NAT in FortiOS provides information about how NAT works in FortiOS and the combinations of NAT that you can use in your configuration. This section explains how the different modes, such as Transparent mode, work and how the FortiGate unit behaves when in each of these modes.

Firewall components provides in-depth information about the firewall components that help in creating a FortiGate firewall configuration.

Security policies explains what security policies are, as well as how these rules work to help protect your network. This section also explains the importance of how security policies are ordered within the security policy list, and describes the different policies that can be created for different firewall configurations.

Monitoring firewall traffic explains how you can monitor traffic within the web-based manager using the Session and Policy Monitoring pages.

Internet Protocol version 6 (IPv6) explains how IPv6 can be implemented in FortiOS, as well as what features support IPv6, such as IPsec VPN and dynamic routing. This section also explains a high-level summary of IPv6.

Advanced FortiGate firewall concepts explains the advanced firewall features that you may want to configure for your network, as it expands. This section explains advanced firewall features that include stateful inspection of SCTP traffic, port pairing (Transparent mode only), and adding NAT security policies in Transparent mode.

Understanding the FortiGate firewall

The FortiGate firewall is one of the most important features on the FortiGate unit, allowing not only traffic to flow through, but also, with the help of security policies, scan the traffic for vulnerabilities and misuse and abuse. This type of firewall provides flexibility for expansion in a growing network environment.

This section helps to explain the FortiGate firewall and its role in protecting your network. This section also explains the life of a packet, which helps you to understand how the traffic flows through the FortiGate unit and the role the FortiGate firewall plays in the life of a packet.

The following topics are included in this section:

- What is the FortiGate firewall?
- FortiGate firewall components
- Understanding how a packet travels through the FortiGate unit

What is the FortiGate firewall?

A firewall is, in the simplest of terms, a device that permits or denies network traffic based on a set of rules. For the FortiGate firewall, it can do this and much more. The FortiGate firewall scans the network traffic, and based on the set of rules (in Fortinet, however, these rules are called security policies), determines what action needs to be taken. The action may be to quarantine a virus that the FortiGate unit finds, or to record the activity, or both. These security policies provide the information the FortiGate unit needs to determine what to do with the incoming and outgoing traffic.

At the heart of these networking security functions, is the security policies. Security policies control all traffic attempting to pass through the FortiGate unit, and between FortiGate interfaces, zones, and VLAN subinterfaces. They are instructions the FortiGate unit uses to decide connection acceptance and packet processing for traffic attempting to pass through. When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number), and attempts to locate a security policy matching the packet.

Security policies can contain many instructions for the FortiGate unit to follow when it receives matching packets. Some instructions are required, such as whether to drop or accept and process the packets, while other instructions, such as logging and authentication, are optional. It is through these policies that the FortiGate unit permits or denies the packets to pass through to the network, who gets priority (bandwidth) over other users, and when the packets can come through.

FortiGate firewall components

The FortiGate firewall is comprised of many different features that provides flexibility for the specific needs of your network, both now and as it grows. These features are:

- interfaces (including VLANs)
- zones
- unified threat management (UTM)
- firewall addresses (this includes IPv4 and IPv6, IP pools,. wildcard addresses and netmasks, and geography-based addresses)
- monitoring traffic
- traffic shaping and per-ip traffic shaping (advanced)
- firewall schedules
- services (such as AOL, DHCP and FTP)
- logging traffic (advanced)
- QoS (advanced)
- identity-based policies (advanced)
- endpoint security (advanced)

All of these components each provide an important role in configuring your FortiGate firewall. For example, the administrator applies the PING admin access to the wan1 interface so that he or she can ping this external interface and verify that Internet traffic is hitting the internal to wan1 security policy. If there was no PING admin access applied to the external interface, the administrator could not properly verify if traffic is hitting the policy.

For more in-depth explanations of these components, see the "Firewall components" on page 23.

How the firewall components create a FortiGate firewall and help in protecting your network

The firewall components each help in protecting your network, as well as helping traffic to flow better through the network, for example traffic shaping helps to load balance traffic on your network.

The following explains how all of the firewall components get combined to create the FortiGate firewall.

1. **In *System* > *Network* > *Interface*, create VLAN subinterfaces for each department: sales, marketing and engineering.**

   These VLAN subinterfaces will be grouped into a zone and the zone will then be applied to a security policy.

2. **Create a zone for the VLAN subinterfaces.**
3. **In *Firewall Objects* > *Address* > *Address*, create the IP address ranges that are required: one for sales, one for marketing, and one for engineering.**

   Each of these ranges corresponds to the departments that have these IP address ranges. For example, sales has 172.16.120.100 - 172.16.120.200.

4. **Create a firewall schedule that allows sales and marketing Internet access all day; create another firewall schedule that allows engineering access to the Internet only during their lunch break.**

   By creating two different firewall schedules, you can block access for one group for a specified time period, and allow another group all day access.

5. **Group the firewall schedules together so that you can apply them both to a security policy.**

6 **Create a virtual IP address that will be used to allow Internet users access to a web server on your DMZ network.**

7 **In *Policy* > *Policy* > *Policy*, create the following:**

- a security policy that allows Internet users access to the web server
- a security policy that applies the firewall schedule group for Internet access for the sales, marketing and engineering departments (this applies the zone)
- a deny policy that blocks FTP downloads

8 **With all the policies now in the list, arrange them so that the most important policies are first, and least important are last. The list order is:**

- deny policy
- security policy that allows Internet users access to the web server
- security policy for sales, engineering and marketing that allows Internet access

Now that all the policies are in the correct order, you need to test that all are working properly.

9 **To verify that traffic is hitting the policies, verify that there is a packet count increase occurring in the Count column of each of the policies in the policy list. Troubleshoot any issues using the diagnose sniffer and diagnose debug flow commands in the CLI.**

By testing that traffic is hitting the policies that you just created, you can see whether you need to solve any issues or not. When you use the `diagnose` commands, you can see detailed information about the traffic hitting the policy.

10 **Back up the configuration after testing and troubleshooting.**

By backing up the changes your made to the configuration, you ensure that a current configuration of this FortiGate firewall configuration is available at any time.

Understanding how a packet travels through the FortiGate unit

Directed by security policies, a FortiGate unit screens network traffic from the IP layer up through the application layer of the TCP/IP stack. The FortiGate firewall plays an important role in how the packet travels through the FortiGate unit out to its destination. The following explains how the packet travels through the FortiGate unit and how the FortiGate firewall plays a role in the life of a packet.

The FortiGate unit performs three types of security inspection:

- stateful inspection, that provides individual packet-based security within a basic session state
- flow-based inspection, that buffers packets and uses pattern matching to identify security threats
- proxy-based inspection, that reconstructs content passing through the FortiGate unit and inspects the content for security threats.

Each inspection component plays a role in the processing of a packet as it traverses the FortiGate unit en route to its destination. When you understand these inspections, you will understand the packet's journey through the FortiGate unit and how the FortiGate firewall helps the packet along to its destination.

For more information about how packets travel through the FortiGate unit, see the Troubleshooting chapter in the *FortiOS Handbook*. The following explains, in a high-level description, of how a packet travels through the FortiGate unit.

How packets flow in and out of the FortiGate unit

The following provides a high-level description of the steps a packet takes when it enters the FortiGate unit, travelling to its destination, the internal network. Similar steps occur for outbound traffic; they are just in reverse.

1. **An incoming packet enters the external interface of the FortiGate unit to start its journey through to the internal network. This is called ingress. During ingress, the following processes occur:**
   - DoS Sensor
   - IP integrity header checking
   - IPsec
   - Destination NAT (DNAT)
   - Routing
2. **After the Routing process finishes, the stateful inspection engine processes the packet, and does the following:**
   - Session Helpers
   - Management Traffic
   - SSL VPN
   - User Authentication
   - Traffic Shaping
   - Session Tracking
   - Policy lookup
3. **If nothing comes from the stateful inspection engine, then the packet travels to the UTM scanning process. This process may have either a flow-based or proxy-based inspection engine that also processes the packet.**
4. **If nothing matches the UTM rules, the packet then travels to other processing steps, which include:**
   - IPsec
   - NAT (Source NAT)
   - Routing
   - Internal Interface
5. **After step 4 is finished, the packet travels out of the internal interface of the FortiGate unit, heading towards its final destination, the internal network. This is referred to as Egress.**

Figure 1: Packet flow

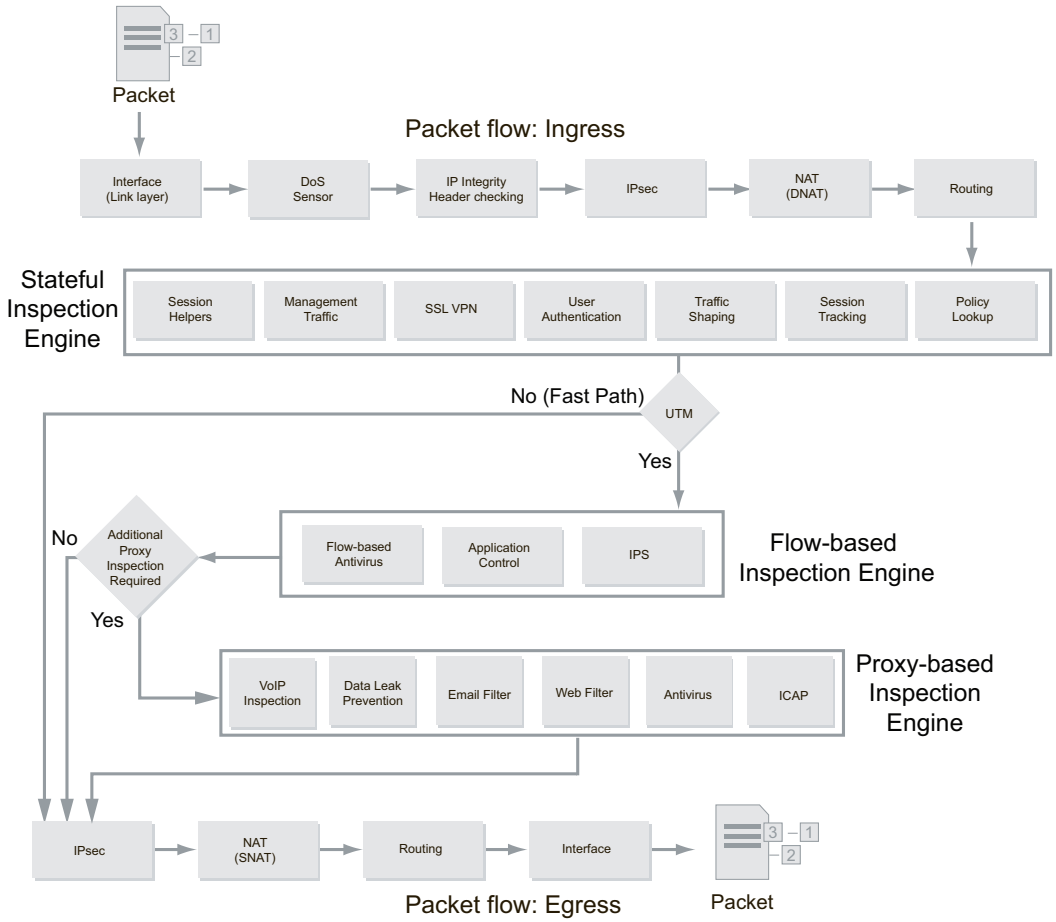

Working with NAT in FortiOS

This section explains NAT and the NAT/Route mode of the FortiGate unit, as well as Transparent mode and its role with NAT. This section also explains the types of NAT that FortiOS supports, including combinations of NAT that you can configure in FortiOS.

This section also includes information about Route mode and how it behaves in FortiOS.

The following topics are included in this section:

- NAT in FortiOS
- Types of NAT in FortiOS
- Combining types of NAT

NAT in FortiOS

Network address translation (NAT) translates one IP address (either a source IP address or destination IP address) for another IP address. NAT in FortiOS, however, can translate IP addresses in many different ways, providing the flexibility you need for your specific network requirements. For example, you can use the Central NAT table to help in translating multiple IP addresses.

When configuring NAT in FortiOS, you should also know how it works within the different modes that the FortiGate unit can be configured in.

This topic contains the following:

- NAT/Route mode
- Route mode
- Transparent mode

NAT/Route mode

In NAT/Route mode, the FortiGate unit is visible to the network that is connected to. All of its interfaces are on different subnets. Each interface it is connected to a network that must be configured with an IP address that is valid for that subnetwork.

NAT/Route mode is typically used when the FortiGate unit is deployed as a gateway between private and public networks. In its default NAT mode configuration, the FortiGate unit functions as a firewall. Security policies control communications through the FortiGate unit to both the Internet and between internal networks. In NAT/Route mode, the FortiGate unit performs network address translation before IP packets are sent to the destination network. For example, a company has a FortiGate unit as their interface to the Internet. The FortiGate unit also acts as a router to multiple sub-networks within the company.

In Figure 2, the FortiGate unit is set to NAT/Route mode and is connected to a network. By using this mode, the FortiGate unit can have a designated port for the Internet, and the internal segments are behind the FortiGte unit, which are invisible to the public access. The FortiGate unit translates IP addresses passing through it to route the traffic to the correct subnet on the Internet.

Figure 2: An example of a FortiGate unit in NAT/Route mode on a network

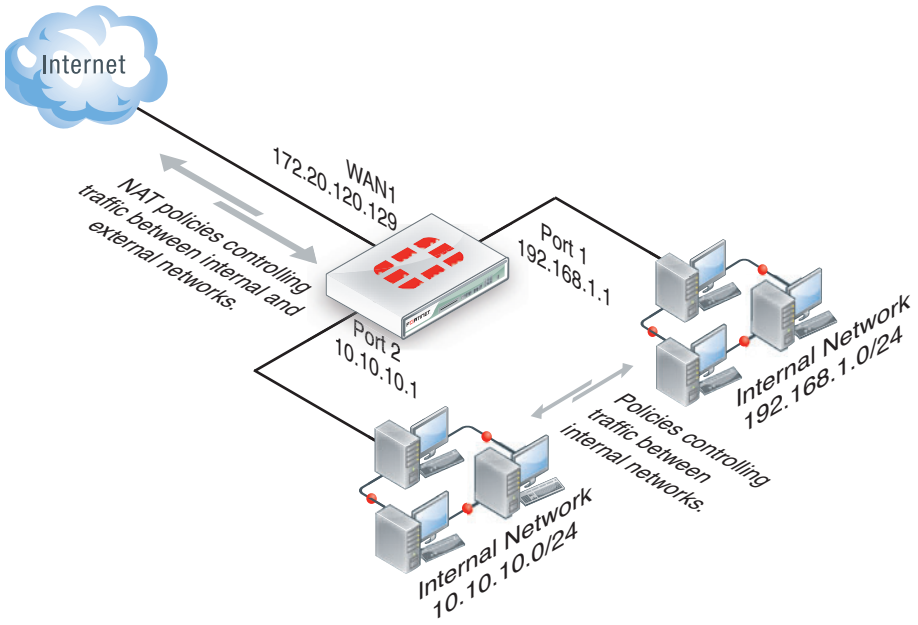

Route mode

In Route mode, the FortiGate unit is only routing traffic, not translating the IP addresses. In this mode, the FortiGate unit acts similar to a switch, passing the packet along to the destination network. This mode is not to be confused with Transparent mode, which is invisible on the network; rather, in Route mode, the FortiGate unit is visible to the network, but does only routing.

The FortiGate unit is used in Route mode whenever no NAT translation needs to be done. For example, you want to connect two separate subnets without using NAT.

You must select NAT/Route mode when configuring the FortiGate unit for Route mode.

Figure 3: An example of a FortiGate unit in Route mode on a network

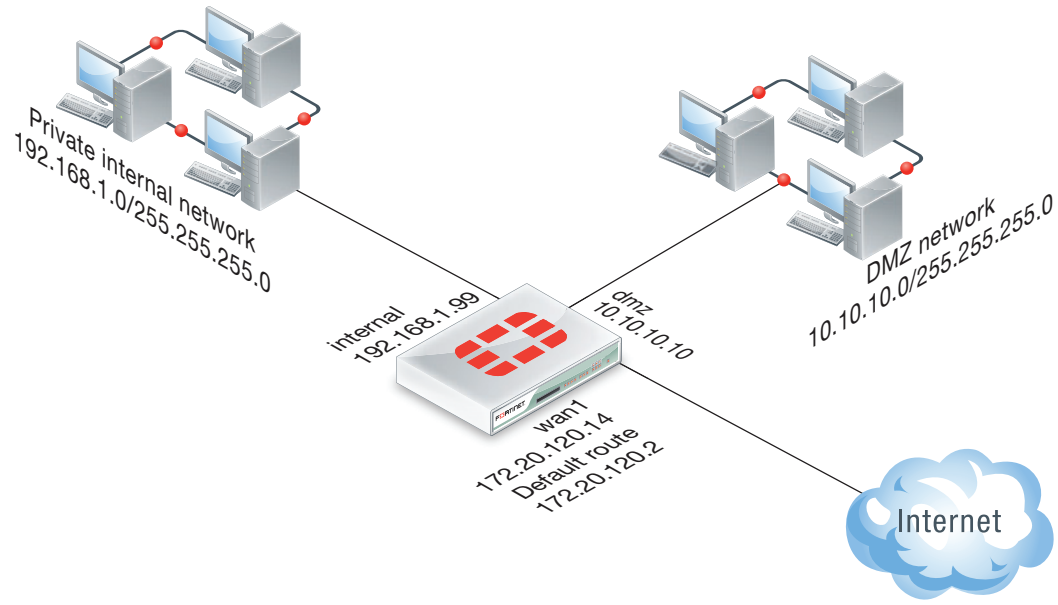

Transparent mode

In Transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. If you want to configure the FortiGate unit in Transparent mode, all you need to do is to configure a management IP address and a default route.

You would typically use Transparent mode on a private network behind an existing firewall or behind a router. In Transparent mode, the FortiGate unit functions as a firewall and can even perform NAT. Security policies control communications through the FortiGate unit to the Internet and internal network. Traffic cannot pass through until you add security policies when the FortiGate unit is in Transparent mode.

In Transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools. If you want NAT to be performed in Transparent mode, you must configure two management IP addresses that are on different subnets.

Figure 4: A FortiGate unit in Transparent mode

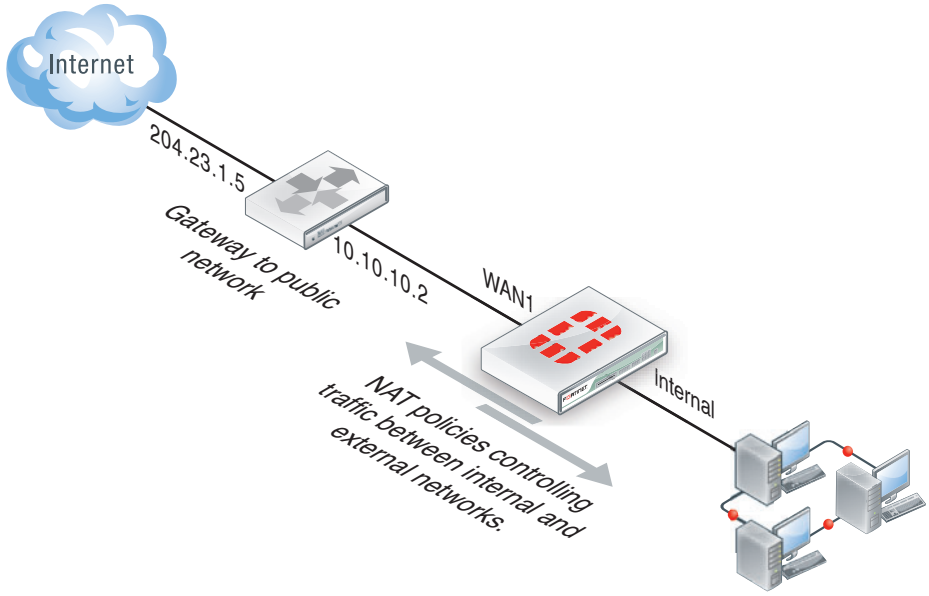

Types of NAT in FortiOS

There are many types of NAT that are available, some you may already know such as port address translation (PAT). The following explains these types of NAT that are available in FortiOS.

This topic contains the following:

- Static NAT (SNAT)
- Dynamic NAT (DNAT)

Static NAT (SNAT)

Static NAT, or source address translation (SNAT), is when a static source IP address is translated by NAT to another source IP address. In FortiOS, when a packet with a specific source address is accepted by a security policy with NAT enabled, the source address is swapped with another IP address. For example, you want to allow a web server on a private network that is protected by a FortiGate unit to connect to the Internet; the web server has a static IP address of 10.10.30.10 and the external interface of the unit is 172.20.120.233; when the packet is received at the FortiGate's internal interface, it is translated from 10.10.30.10 to 172.20.120.133, and forwards the packet out to the Internet.

Static NAT is used when configuring basic security policies. For example, you want users on a private network to connect to the Internet.

When configuring static NAT security policies, there are several steps that must be configured prior to configuring the actual security policy. For example, for static DNAT, you must configure a virtual IP address that maps to a specific destination address.

Static Destination NAT (SDNAT)

As stated for static NAT, the same is true for static destination address translation, or SDNAT, whereby a packet with a specific destination address is accepted by a security policy with NAT enabled, the destination address is swapped with another destination address.

Static NAT port forwarding

There is also static port forwarding, which acts similarly to static DNAT, translating a destination address and port number to another destination address and port number. The difference is that port forwarding requires a virtual IP address so that the FortiGate unit can properly translate the port number.

When a packet with a destination address to be translated is accepted by a security policy (with DNAT enabled), and a virtual IP with an external port mapped to that address's port, then the FortiGate unit swaps the packet's destination address with the other IP address, and its port number with the external port.

Dynamic NAT (DNAT)

As subnets grow larger, more work is required to set network address translation with each additional client. Rather than assigning static addresses, an administrator may want to set up IP pools. IP pools are ranges of addresses that clients on a subnet can use to send and receive packets, as well as which FortiGate units can use to translate the addresses of packets going through them. This type of translation is known as dynamic NAT, when address translation is done on a flexible or "many-to-one" basis using IP pools.

IP pools do not randomly assign addresses, rather, each IP pool is a prioritized list of IP addresses. When a client is assigned an IP address from the IP pool, it retains that address. Another client that requires an address is then assigned the next IP address from that IP pool list. When the range of virtual IPs are used instead of IP pools, these virtual IPs are prioritized in the same type of list.

Dynamic source address translation

Dynamic source address translation has economies of scale for larger subnets and more flexible subnets, enabling network infrastructure to change without the hassle of reconfiguring addresses after every change. Dynamic source address NAT or DNAT translates many source addresses as defined by an IP pool. Whenever a packet with the specific source address to be translated is accepted by a security policy with source NAT enabled, the FortiGate unit swaps the packet's source address with the other IP address selected from the IP pool.

For example, an organization may want packets leaving the FortiGate unit for the Internet to have source Ips in the range of 172.16.0.1-10. This means that packets accepted by a firewall policy must have their source addresses translated to an address in this range before being forwarded to the Internet. So if the server on the private network with the address 10.0.0.1 has its source IP translated to 172.16.0.1, then the next available source IP in the IP pool will be 172.16.0.2, which a server with address 10.0.0.2 can use.

Dynamic destination address

Dynamic destination address NAT (or DDNAT) translates one range of destination addresses to another range of destination addresses. Whenever a packet within the specified range of destination addresses to be translated is accepted by a security policy with DNAT is enabled, the FortiGate unit swaps the packet's destination address with one of the addresses from the other specified range.

For example, to allow customers from the Internet to connect to several web servers protected by a FortiGate unit, you require a range of Internet addresses (for example, 172.16.0.1-10), enough for each protected web server, and a range of real addresses (for example, 10.0.0.1-10) for each web server. When a packet is received at the external interface of the FortiGate unit with a destination IP address within the Internet range of addresses, the FortiGate unit translates the destination address of the packet to the real address and forwards the packet to the web server on the network protected by the FortiGate unit.

Dynamic port forwarding

Dynamic port forwarding translates one range of destination addresses and ports to another range of destination addresses and ports. Whenever a packet with a specified destination address to be translated is accepted by a security policy with destination NAT enabled, and a virtual IP with an external port mapped to that address's port, then the FortiGate unit swaps the packet's destination address with the other IP address, and its port number with the external port.

For example, to allow customers from the Internet to connect to web servers protected by a FortiGate unit, you require a range of Internet addresses (for example, 172.16.0.1-10) and a range of port numbers (for example, 80-89), and a range of ports numbers to be mapped to (for example, 8080-8089). When a packet is received at the external interface of the FortiGate unit with the 172.16.0.3 destination IP address and port number 8082, the FortiGate unit translates that address to 10.0.0.3 and port number to 82, and then forwards the packet to the web server.

Combining types of NAT

In FortiOS, you can combine a number of NAT features to get the best firewall configuration possible for your network requirements. NAT combinations include Double NAT, which is combining IP pool with virtual IP, and using VIP range for SNAT and static one-to-one mapping.

These combinations can help you when creating your FortiGate firewall configuration. The combinations help when you have multiple addresses (IP pools) and when you need to use a virtual IP address with the IP pool. An example of this combination is called Double NAT.

You can also combine dynamic NAT types, such as dynamic source address translation, to help you with creating the FortiGate firewall using dynamic NAT. An example of this combination is using the Central NAT table.

When considering your FortiGate firewall configuration, you should also consider how to combine NAT types. By combining NAT types, you can easily use multiple addresses when configuring security policies, as well as when you want to provide specific NAT translations, such as using dynamic source NAT that will not change the source port; this combination allows for the handling of specific protocols or services that function only if they use a specific port and that port does not change.

The following are some combinations of NAT that you can use in your FortiGate firewall configuration:

- Double NAT
- Central NAT table (similar to IP pools)
- virtual IP range for SNAT
- static one-to-one mapping
- dynamic source NAT (also known as one-to-one source NAT)
- dynamic source NAT (this uses Dynamic IP pool and a virtual IP)

Firewall components

The FortiGate unit's primary purpose is to act as a firewall to protect your networks from unwanted attacks and to control the flow of network traffic. The firewall consists of many different and important components so that you can better protect your network as your network requirements grow. This section explains these components.

The following topics are included in this section:

- Using Interfaces and zones in the FortiGate firewall
- Understanding the firewall address component
- UTM profiles

Using Interfaces and zones in the FortiGate firewall

Interfaces and zones are used when configuring security policies to define incoming and outgoing traffic. For example, in an internal to wan 1 security policy, the internal interface is where traffic is coming in, and the wan 1 interface is where the traffic is going out to. When the FortiGate unit sees that traffic came in using the internal interface, and needs to leave using the wan 1 (or external interface), the security policy internal to wan1 is matched to the traffic and additional rules are applied to the traffic as well.

Interfaces, either virtual or physical, can be applied to security policies. VLAN subinterfaces are virtual interfaces that can be applied to security policies to control and direct traffic on those subinterfaces. VLAN subinterfaces are interfaces that are part of one of the main interfaces, for example, wan1. For more information about VLAN subinterfaces and how to configure them, see the System Admin chapter of the *FortiOS Handbook*.

Zones provide the option of grouping multiple FortiGate interfaces, both virtual and physical, that you can then apply to security policies to control the incoming and outgoing traffic on those interfaces. By using zones, you can easily group multiple interfaces and VLAN subinterfaces together to help simplify creating security policies where a number of network segments can use the same policy and UTM settings.

How to apply VLANs and zones and to a security policy

The following explains how to create three VLAN subinterfaces, grouping these subinterfaces into a zone, and then applying the zone to a security policy. The security policy will control the traffic for these VLAN subinterfaces.

1. **Create three VLANs in *System* > *Network* > *Interface* for engineering, sales and marketing on the internal interface.**

   These three VLANs will be grouped together to create a zone which will then be applied to the security policy. The zone will be applied to the policy instead of the individual VLANs.
2. **Group the VLANs into a zone.**
3. **Create DHCP servers for each of the VLAN subinterfaces in *System* > *Network* > *DHCP*.**

4 **Create the security policy for the zone to control traffic in *Policy* > *Policy* > *Policy*.**

   In the *Source Interface/Zone* list, you would instead choose the zone. The *Destination Interface/Zone* is the external interface, wan1. By choosing the zone, you apply all the subinterfaces at once.

5 **Select *Enable NAT* and *Use Destination Interface Address*; ensure that *Log Allowed Traffic* is also enabled so that you can use the logs to help determine if traffic is hitting the security policy.**

Understanding the firewall address component

Firewall addresses in FortiOS provide flexibility when configuring access control over the network traffic. When this document talks about firewall addresses, this encompasses:

- IP addresses and netmasks
- IP pools (this can include the Central NAT table)
- virtual IP addresses
- geography-based addresses
- IPv4 addresses
- wildcard addresses and netmasks
- Fully Qualified Domain Name addresses (FQDN)
- IP address groups

Firewall addresses help define the network addresses that you use when configuring a security policy's source and destination address. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

A firewall address can contain one or more network addresses. Network addresses can be represented by an IP address with a netmask, an IP address range, or a fully qualified domain name (FQDN).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- a single computer, such as `192.45.46.45`
- a subnetwork, such as `192.168.1.0` for a class C subnet
- `0.0.0.0`, which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- netmask for a single computer: `255.255.255.255`, or `/32`
- netmask for a class A subnet: `255.0.0.0`, or `/8`
- netmask for a class B subnet: `255.255.0.0`, or `/16`
- netmask for a class C subnet: `255.255.255.0`, or `/24`
- netmask including all IP addresses: `0.0.0.0`

Valid IP address and netmask formats include:

- x.x.x.x/x.x.x.x, such as `192.168.1.0/255.255.255.0`
- x.x.x.x/x, such as `192.168.1.0/24`

An IP address of 0.0.0.0 with a netmask 255.255.255.255 is not a valid firewall address.

When representing hosts by an IP address range, the range indicates hosts with continuous IP addresses in a subnet, such as `192.168.1.[2-10]`, or `192.168.1.*` to indicate the complete range of hosts on that subnet. Valid IP Range formats include:

- x.x.x.x-x.x.x.x, such as `192.168.110.100-192.168.110.120`
- x.x.x.[x-x], such as `192.168.110.[100-120]`
- x.x.x.*, such as `192.168.110.*`

When representing hosts by an FQDN, the domain name can be a subdomain, such as mail.example.com. A single FQDN firewall address may be used to apply a security policy to multiple hosts, as in load balancing and high availability (HA) configurations. FortiGate units automatically resolve and maintain a record of all addresses to which the FQDN resolves. Valid FQDN formats include:

- <host_name>.<second_level_domain_name>.<top_level_domain_name>, such as `mail.example.com`
- <host_name>.<top_level_domain_name>

Be cautious when employing FQDN firewall addresses. By using a fully qualified domain name in a security policy, while convenient, does present some security risks, because policy matching then relies on a trusted DNS server. If the DNS server should ever be compromised, security policies requiring domain name resolution may no longer function properly.

This topic contains the following:

- IP addresses for self-originated traffic
- IP pools
- IP Pools for security policies that use fixed ports
- Source IP address and IP pool address matching
- Geography-based addressing
- Wildcard addresses
- Fully Qualified Domain Name addresses
- Address groups
- Virtual IP addresses

IP addresses for self-originated traffic

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email

- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Configuration of these services is performed in the CLI. In each instance, there is a command set `source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
  set ntpsyn enable
  set syncinterval 5
  set source-ip 192.168.4.5
end
```

To see which services are configured with source-ip settings, use the `get` command:

```
get system source-ip status
```

The output will appear similar to the sample below:

```
NTP: x.x.x.x
DNS: x.x.x.x
SNMP: x.x.x.x
Central Management: x.x.x.x
FortiGuard Updates (AV/IPS): x.x.x.x
FortiGuard Queries (WebFilter/SpamFilter): x.x.x.x
```

IP pools

An IP pool defines a single IP address or a range of IP addresses. A single IP address in an IP pool becomes a range of one IP address. For example, if you enter an IP pool as 1.1.1.1, the IP pool is actually the address range, 1.1.1.1 to 1.1.1.1. Use IP pools to add NAT policies that translate source addresses to addresses randomly selected from the IP pool, rather than the IP address assigned to that FortiGate interface. You can use the Central NAT table as a way to configure IP pools. For more information, see “Central NAT table” on page 65.

If a FortiGate interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools.

For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

- (1.1.1.0-1.1.1.255) and (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20

The port2 interface overlap IP range with IP_pool_2 is:

- (2.2.2.0-2.2.2.255) & (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20

The port2 interface overlap IP range with IP_pool_3 is:

- (2.2.2.0-2.2.2.255) & (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select *Enable NAT* in a security policy and then select *Dynamic IP Pool.* Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool.

IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

IP Pools for security policies that use fixed ports

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable `fixedport` when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
  edit policy_name
    ...
    set fixedport enable
    ...
end
```

However, enabling `fixedport` means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select *Dynamic IP pool* in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Source IP address and IP pool address matching

When the source addresses are translated to the IP pool addresses, one of the following three cases may occur:

Scenario 1: The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable `fixedport` in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Original address	Change to
192.168.1.1	172.16.30.1
192.168.1.2	172.16.30.2
.....	.....
192.168.1.254	172.16.30.254

Scenario 2: The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism.

If you enable `fixedport` in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Original address	Change to
192.168.1.1	172.16.30.10
192.168.1.2	172.16.30.11
.....	.....
192.168.1.10	172.16.30.19
192.168.1.11	172.16.30.10
192.168.1.12	172.16.30.11
192.168.1.13	172.16.30.12
.....	.....

Scenario 3: The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

Original address	Change to
192.168.1.1	172.16.30.10
192.168.1.2	172.16.30.11
192.168.1.3	172.16.30.12
No more source addresses	172.16.30.13 and other addresses are not used

Geography-based addressing

An option is available to add a geography-based address scheme. With this type of addressing, you indicate the geographic region, or country. The FortiGate unit includes an internal list of countries and IP addresses based on historical data from the FortiGuard network.

IPv6 does not support geography-based addressing. This feature is for IPv4 addresses only.

When used in security policies, traffic originating or going to a particular country can be logged, blocked or specific filtering applied.

In the following examples, an geographic-based address for China is added for the WAN1 port.

To add a geography-based address - web-based manager

1. Go to *Firewall Objects* > *Address* > *Address* and select *Create New*.
2. Enter the *Name* of `China`
3. For the *Type*, select *Geography*.

4 From the *Country* list, select *China*.

5 Select the *Interface* of *WAN1*.

6 Select *OK*.

To add a geography-based address - CLI

```
config firewall address
  edit China
    set type geography
    set country CN
    set interface wan1
end
```

You can use a `diagnose` command to view more information about geography-based addressing. The command displays country and address information for the countries that have been added to firewall addresses.

```
diagnose firewall ipgeo {country-list | ip-list | ip2country}
```

Where:

- `country-list` shows all of the countries that have been added to a firewall address.
- `ip-list` shows the IP addresses of a specified country or all of the countries added to firewall addresses.
- `ip2country` shows the country of origin for a specified IP address. The address must be assigned to one of the countries that has been added to a firewall address.

Wildcard addresses

Wildcard addresses are addresses that identify ranges of IP addresses, reducing the amount of firewall addresses and security policies required to match some of the traffic on your network. Wildcard addresses are an advanced feature, usually required only for complex networks with complex firewall filtering requirements. By using these wildcard addresses in the firewall configuration, administrators can eliminate creating multiple, separate IP addresses and then grouping them to then apply to multiple security policies.

A wildcard address consists of an IP address and a wildcard netmask, for example, 192.168.0.56 255.255.0.255. In this example, the IP address is 192.168.0.56 and the wildcard netmask is 255.255.0.255. The IP address defines the networks to match and the wildcard netmask defines the specific addresses to match on these networks.

In a wildcard netmask, zero means ignore the value of the octet in the IP address, which means the wildcard firewall address matches any number in this address octet. This also means that the number included in this octet of IP address is ignored and can be any number. Usually, if the octet in the wildcard netmask is zero, the corresponding octet in the IP address is also zero.

In a wildcard netmask, a number means match addresses according to how the numbers translate into binary addresses. For example, the wildcard netmask is 255; the wildcard address will only match addresses with the value for this octet that is in the IP address part of the wildcard address. So, if the first octet of the IP address is 192 and the first octet of the wildcard netmask is 255, the wildcard address will only match addresses with 192 in the first octet.

In the above example, the wildcard address 192.168.0.56 255.255.0.255 would match the following IP addresses:

192.168.0.56, 192.168.1.56, 192.168.2.56, ..., 192.168.255.56

The wildcard addresses 192.168.0.56 255.255.0.255 and 192.168.1.56 255.255.0.255 define the same thing since the 0 in the wildcard mask means to match any address in the third octet.

If we use the wildcard address 172.0.20.10 255.0.255.255, it would match the following IP addresses:

172.1.20.10, 172.2.20.10, 172.3.20.10, ..., 172.255.20.10

In a wildcard netmask, a number other than 255 matches multiple addresses for this octet. You can perform a binary conversion to calculate the addresses that would be matched by a given value. For example, to create the IP address and wildcard netmask to match the following network addresses:

```
192.168.32.0/24
192.168.33.0/24
192.168.34.0/24
192.168.35.0/24
192.168.36.0/24
192.168.37.0/24
192.168.38.0/24
192.168.39.0/24
```

Table 1 shows how to write the third octet for these networks according to the octet bit position and address value for each bit.

Table 1: Octet bit position and address value for each bit

Decimal	128	64	32	16	8	4	2	1
32	0	0	1	0	0	0	0	0
33	0	0	1	0	0	0	0	1
34	0	0	1	0	0	0	1	0
35	0	0	1	0	0	0	1	1
36	0	0	1	0	0	1	0	0
37	0	0	1	0	0	1	0	1
38	0	0	1	0	0	1	1	0
39	0	0	1	0	0	1	1	1
	M	M	M	M	M	D	D	D

Since the first five bits match, the networks can be summarized into one network (192.168.32.0/21 or 192.168.32.0 255.255.248.0). All eight possible combinations of the three low-order bits are relevant for the network ranges. The wildcard address that would match all of these subnet addresses can be written as 192.168.32.0 255.255.248.0.

Wildcard addresses are similar to routing access list wildcard masks. You add routing access lists containing wildcard masks using the `config router access-list` command. However, router access list wildcard masks use the inverse of the masking system used for firewall wildcard addresses. For the router access list wildcard masks, zero (0)means match all IP addresses and one (1)means ignore all IP addresses. So to match IP addresses 192.168.0.56, 192.268.1.56, 192.168.2.56, ... 192.168.255.56 you would use the following router access IP address prefix and wildcard mask: 192.168.0.56 0.0.255.0.

Wildcard firewall addresses are configured only in the CLI. The following is an example of how to configure a wildcard firewall address.

```
config firewall address
```

```
edit example_wildcard_address
  set type wildcard
  set wildcard 192.168.0.56 255.255.0.255
end
```

Using wildcard addresses in the firewall configuration

The following example shows how wildcard addresses can be applied to network traffic. This example consists of a security policy where both the source and destination addresses are firewall wildcard addresses.

Source Address: 10.129.5.0 255.127.7.0

Destination Address: 10.129.0.10 255.127.7.255

A security policy with these source and destination addresses would permit:

- A device with IP address 10.129.5.100 to connect through the FortiGate unit to IP address 10.129.0.10
- A device with IP address 10.129.13.100 to connect through the FortiGate unit to IP address 10.129.8.10
- A device with IP address 10.129.21.100 to connect through the FortiGate unit to IP address 10.129.0.10

In another example of wildcard addresses, the following shows how only odd numbered addresses get allowed through:

1. Create wildcard address 4.2.2.0/255.255.255.1.

   This is configured in the CLI.

2. Create a deny security policy that uses the wildcard address, 4.2.2.0.

   The results are that only the odd-numbered 4.2.2.0 addresses are allowed in; all other addresses are blocked.

Fully Qualified Domain Name addresses

Be cautious when employing FQDN firewall addresses. Using a fully qualified domain name in a security policy, while convenient, does present some security risks, because policy matching then relies on a trusted DNS server. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Using Fully Qualified Domain Name (FQDN) addresses in security policies has the advantage of causing the FortiGate unit to keep track of DNS TTLs and adapt as records change. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache. The FortiGate unit will query the DNS for an amount of time specified, in seconds, and update the cache as required. This feature can reduce maintenance requirements for changing firewall addresses for dynamic IP addresses. This also means that you can create security policies for networks configured with dynamic addresses using DHCP.

You specify the TTL time in the CLI only. For example, to set the TTL for 30 minutes on an FQDN of www.example.com on port 1, enter the following commands:

```
config firewall address
  edit FQDN_example
    set type fdqn
    set associated-interface port 1
    set fqdn www.example.com
```

```
        set cache-ttl 1800
    end
```

Address groups

Similar to zones, if you have a number of addresses or address ranges that require the same security policies, you can put them into address groups, rather than creating multiple similar policies. Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to *Any* — addresses whose selected interface is *Any* are bound to a network interface during creation of a security policy, rather than during creation of the firewall address.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of *Any*, they can be grouped, even if the addresses involve different networks.

You cannot mix IPv4 firewall addresses and IPv6 firewall addresses in the same address group.

Virtual IP addresses

In FortiOS, virtual IP addresses (VIPs) can be used when configuring security policies to translate IP addresses and ports of packets received by a network interface. When the FortiGate unit receives inbound packets matching a security policy whose *Destination Address* field is a virtual IP, the FortiGate unit applies NAT, replacing packets's IP addresses with the virtual IP's mapped IP address.

VIPs can specify translation of packets' port numbers and/or IP addresses for both inbound and outbound connections. In Transparent mode, virtual IPs are available only in the CLI.

VIP addresses are typically used to map external (public) to internal (private) IP addresses for Destination NAT (DNAT).

Grouping virtual IPs

You can organize multiple virtual IPs into a virtual IP group to simplify your security policy list. For example, instead of having five identical policies for five different but related virtual IPs located on the same network interface, you might combine the five virtual IPs into a single virtual IP group, which is used by a single security policy.

Security policies using VIP Groups are matched by comparing both the member VIP IP addresses) and port numbers).

Match-vip

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. Use the following command syntax to enable this feature. By default, it is disabled.

```
config firewall policy
```

```
edit <vip_policy_name>
  set match-vip {disable | enable}
end
```

How to use match-vip

In this example, a deny security policy has already been configured that blocks FTP sessions. A virtual IP address will be configured in this example and then applied to a security policy that allows Internet users access to a web server on the company's DMZ network.

1 **Create the virtual IP address in *Firewall Objects* > *Virtual IP* > *Virtual IP*.**

   This address is called vip-dmz. You can configure the virtual IP address solely in the CLI. This would eliminate having to go back and forth.

2 **Log in to the CLI and enter the following commands:**

```
config firewall policy
  edit vip-dmz
    set match-vip enable
  end
```

3 **Create the virtual IP security policy.**

   For this security policy, you need to turn on logging within the security policy.

4 **Test the policy to view the activity that is occurring with the `match-vip` command enabled.**

Services

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined in firewall services and protocols on the FortiGate unit. These predefined services and protocols are defaults, and cannot be edited or removed. However, if you require different services, you can create custom services.

To view the predefined servers, go to *Firewall Objects* > *Service* > *Predefined*.

If there is a service that does not appear on the list, or you have a unique service or situation, you can create your own custom service. You need to know the ports, IP addresses or protocols of that particular service or application uses, to create the custom service.

Predefined service list

Many well-known traffic types have been predefined in firewall services. These predefined services are defaults, and cannot be edited or removed. However, if you require different services, you can create custom services.

Predefined services are located in *Firewall Objects* > *Service* > *Predefined*. Table 2 lists the Fortinet firewall predefined services.

Table 2: Predefined services

Service name	Description	Protocol	Port
AFS3	Advanced File Security Encrypted File, version 3, of the AFS distributed file system protocol.	TCP	7000-7009
		UDP	7000-7009
AH	Authentication Header. AH provides source host authentication and data integrity, but not secrecy. This protocol is used for authentication by IPSec remote gateways set to aggressive mode.	IP	51
ANY	Matches connections using any protocol over IP.	all	all
AOL	America Online Instant Message protocol.	TCP	5190-5194
BGP	Border Gateway Protocol. BGP is an interior/exterior routing protocol.	TCP	179
CVSPSERVER	Concurrent Versions System Proxy Server.CSSPServer is very good for providing anonymous CVS access to a repository.	TCP	2401
		UDP	2401
DCE-RPC	Distributed Computing Environment / Remote Procedure Calls. Applications using DCE-RPC can call procedures from another application without having to know on which host the other application is running.	TCP	135
		UDP	135
DHCP	Dynamic Host Configuration Protocol. DHCP allocates network addresses and delivers configuration parameters from DHCP servers to hosts.	UDP	67 68
DHCP6	Dynamic Host Configuration Protocol for IPv6.	UDP	546, 547
DNS	Domain Name Service. DNS resolves domain names into IP addresses.	TCP	53
		UDP	53
ESP	Encapsulating Security Payload. ESP is used by manual key and AutoIKE IPSec VPN tunnels for communicating encrypted data. AutoIKE VPN tunnels use ESP after establishing the tunnel by IKE.	IP	50
FINGER	A network service providing information about users.	TCP	79
FTP	File Transfer Protocol.	TCP	21
FTP_GET	File Transfer Protocol. FTP GET sessions transfer remote files from an FTP server to an FTP client computer.	TCP	21

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
FTP_PUT	File Transfer Protocol. FTP PUT sessions transfer local files from an FTP client to an FTP server.	TCP	21
GOPHER	Gopher organizes and displays Internet server contents as a hierarchically structured list of files.	TCP	70
GRE	Generic Routing Encapsulation. GRE allows an arbitrary network protocol to be transmitted over any other arbitrary network protocol, by encapsulating the packets of the protocol within GRE packets.	IP	47
GTP (FortiOS Carrier only)	GPRS Tunneling protocol (GTP). GTP is used with GSM and UMTS networks to carry user data within GPRS core networks. FortiOS Carrier can accept and process IPv4 GTP packet.	UDP	2123,2152,3386
H323	H.323 multimedia protocol. H.323 is a standard approved by the International Telecommunication Union (ITU) defining how audiovisual conferencing data can be transmitted across networks. For more information, see the *FortiGate Support for H.323 Technical Note*.	TCP	1720, 1503
		UDP	1719
HTTP	Hypertext Transfer Protocol. HTTP is used to browse web pages on the World Wide Web.	TCP	80
HTTPS	HTTP with secure socket layer (SSL). HTTPS is used for secure communication with web servers.	TCP	443
ICMP_ANY	Internet Control Message Protocol. ICMP allows control messages and error reporting between a host and gateway (Internet).	ICMP	Any
IKE	Internet Key Exchange. IKE obtains authenticated keying material for use with the Internet Security Association and Key Management Protocol (ISAKMP) for IPSEC.	UDP	500, 4500
IMAP	Internet Message Access Protocol. IMAP is used by email clients to retrieves email messages from email servers.	TCP	143
IMAPS	IMAP with SSL. IMAPS is used for secure IMAP communication between email clients and servers. IMAPS is only available on Fortinet units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook.	TCP	993
INFO_ADDRESS	ICMP information request messages.	ICMP	17

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
INFO_REQUEST	ICMP address mask request messages.	ICMP	15
IRC	Internet Relay Chat. IRC allows users to join chat channels.	TCP	6660-6669
Internet-Locator-Service	Internet Locator Service. ILS includes LDAP, User Locator Service, and LDAP over TLS/SSL.	TCP	389
L2TP	Layer 2 Tunneling Protocol. L2TP is a PPP-based tunnel protocol for remote access.	TCP	1701
		UDP	1701
LDAP	Lightweight Directory Access Protocol. LDAP is used to access information directories.	TCP	389
MGCP	Media Gateway Control Protocol. MGCP is used by call agents and media gateways in distributed Voice over IP (VoIP) systems.	UDP	2427, 2727
MMS (FortiOS Carrier only)	MMS tunneling protocol. MMS is used when sending and receiving multimedia content to a mobile phone.	TCP UDP	1755 1024-5000
MS-SQL	Microsoft SQL Server is a relational database management system (RDBMS) produced by Microsoft. Its primary query languages are MS-SQL and T-SQL.	TCP	1433, 1434
MYSQL	MySQL is a relational database management system (RDBMS) which runs as a server providing multi-user access to a number of databases.	TCP	3306
NFS	Network File System. NFS allows network users to mount shared files.	TCP	111, 2049
		UDP	111, 2049
NNTP	Network News Transport Protocol. NNTP is used to post, distribute, and retrieve Usenet messages.	TCP	119
NTP	Network Time Protocol. NTP synchronizes a host's time with a time server.	TCP	123
		UDP	123
NetMeeting	NetMeeting allows users to teleconference using the Internet as the transmission medium.	TCP	1720
ONC-RPC	Open Network Computing Remote Procedure Call. ONC-RPC is a widely deployed remote procedure call system.	TCP	111
		UDP	111
OSPF	Open Shortest Path First. OSPF is a common link state routing protocol.	IP	89
PC-Anywhere	PC-Anywhere is a remote control and file transfer protocol.	TCP	5631
		UDP	5632

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
PING	Ping sends ICMP echo request/replies to test connectivity to other hosts.	ICMP	8
PING6	Ping6 sends ICMPv6 echo request/replies to network hosts to test IPv6 connectivity to other hosts.	ICMP6	58
POP3	Post Office Protocol v3. POP retrieves email messages.	TCP	110
POP3S	Post Office Protocol v3 with secure socket layer (SSL). POP3S is used for secure retrieval of email messages. POP3S is only available on Fortinet units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook.	TCP	995
PPTP	Point-to-Point Tunneling Protocol. PPTP is used to tunnel connections between private network hosts over the Internet. **Note:** Also requires IP protocol 47.		47
		TCP	1723
QUAKE	Quake multi-player computer game traffic.	UDP	26000, 27000, 27910, 27960
RADIUS	Remote Authentication Dial In User Service. RADIUS is a networking protocol that provides centralized access, authorization and accounting management for people or computers to connect and use a network service.	TCP	1812, 1813
RAUDIO	RealAudio multimedia traffic.	UDP	7070
RDP	Remote Desktop Protocol is a multi-channel protocol that allows a user to connect to a networked computer.	TCP	3389
REXEC	Rexec traffic allows specified commands to be executed on a remote host running the rexecd service (daemon).	TCP	512
RIP	Routing Information Protocol. RIP is a common distance vector routing protocol. This service matches RIP v1.	UDP	520
RLOGIN	Remote login traffic.	TCP	513
RSH	Remote Shell traffic allows specified commands to be executed on a remote host running the rshd service (daemon).	TCP	514

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
RTSP	Real Time Streaming Protocol is a protocol for use in streaming media systems which allows a client to remotely control a streaming media server, issuing VCR-like commands such as play and pause, and allowing time-based access to files on a server.	TCP	554, 7070, 8554
		UDP	554
SAMBA	Server Message Block. SMB allows clients to use file and print shares from enabled hosts. This is primarily used for Microsoft Windows hosts, but may be used with operating systems running the Samba daemon.	TCP	139
SCCP	Skinny Client Control Protocol. SCCP is a Cisco proprietary standard for terminal control for use with voice over IP (VoIP).	TCP	2000
SIP	Session Initiation Protocol. SIP allows audiovisual conferencing data to be transmitted across networks. For more information, see the Voice Solutions: SIP chapter of the FortiOS Handbook.	UDP	5060
SIP-MSNmessenger	Session Initiation Protocol used by Microsoft Messenger to initiate an interactive, possibly multimedia session.	TCP	1863
SMTP	Simple Mail Transfer Protocol. SMTP is used for sending email messages between email clients and email servers, and between email servers.	TCP	25
SMTPS	SMTP with SSL. Used for sending email messages between email clients and email servers, and between email servers securely. SMTPS is only available on Fortinet units that support SSL content scanning and inspection. For more information, see the UTM chapter of the *FortiOS Handbook*.	TCP	465
SNMP	Simple Network Management Protocol. SNMP can be used to monitor and manage complex networks.	TCP	161-162
		UDP	161-162
SOCKS	SOCKetS. SOCKS is an Internet protocol that allows client-server applications to transparently use the services of a network firewall.	TCP	1080
		UDP	1080

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
SQUID	A proxy server and web cache daemon that has a wide variety of uses that includes speeding up a web server by caching repeated requests; caching web, DNS and other computer network lookups for a group of people sharing network resources; aiding security by filtering traffic.	TCP	3128
SSH	Secure Shell. SSH allows secure remote management and tunneling.	TCP	22
		UDP	22
SYSLOG	Syslog service for remote logging.	UDP	514
TALK	Talk allows conversations between two or more users.	UDP	517-518
TCP	Matches connections using any TCP port.	TCP	0-65535
TELNET	Allows plain text remote management.	TCP	23
TFTP	Trivial File Transfer Protocol. TFTP is similar to FTP, but without security features such as authentication.	UDP	69
TIMESTAMP	ICMP timestamp request messages.	ICMP	13
TRACEROUTE	A computer network tool used to determine the route taken by packets across an IP network.	TCP	33434
		UDP	33434
UDP	Matches connections using any UDP port.	UDP	0-65535
UUCP	Unix to Unix Copy Protocol. UUCP provides simple file copying.	UDP	540
VDOLIVE	VDO Live streaming multimedia traffic.	TCP	7000-7010
VNC	Virtual Network Computing.VNC is a graphical desktop sharing system which uses the RFB protocol to remotely control another computer.	TCP	5900
WAIS	Wide Area Information Server. WAIS is an Internet search protocol which may be used in conjunction with Gopher.	TCP	210
WINFRAME	WinFrame provides communications between computers running Windows NT, or Citrix WinFrame/MetaFrame.	TCP	1494
WINS	Windows Internet Name Service is Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	TCP	1512
		UDP	1512

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
X-WINDOWS	X Window System (also known as X11) can forward the graphical shell from an X Window server to X Window client.	TCP	6000-6063

Custom service groups

You can organize multiple firewall services into a service group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall services, you might combine the five services into a single address group that is used by a single security policy.

Service groups can contain both predefined and custom services. Service groups cannot contain other service groups.

You can organize multiple firewall services into a service group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall services, you might combine the five services into a single service group that is used by a single security policy.

Service groups can contain both predefined and custom services. Service groups cannot contain other service groups.

Firewall schedules

When you add security policies on a FortiGate unit, those policies are always on, policing the traffic through the device. Firewall schedules control when policies are in effect, that is, when they are on. You can create one-time schedules which are schedules that are in effect only once for the period of time specified in the schedule. You can also create recurring schedules that are in effect repeatedly at specified times of specified days of the week.

You can create a recurring schedule that activates a policy during a specified period of time. For example, you might prevent game playing during office hours by creating a recurring schedule that covers office hours.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next. For example, to prevent game playing except at lunchtime, you might set the start time for a recurring schedule at 1:00 p.m. and the stop time at 12:00 noon. To create a recurring schedule that runs for 24 hours, set the start and stop times to 00.

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall schedules, you might combine the five schedules into a single schedule group that is used by a single security policy.

Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

Schedule groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall schedules, you might combine the five schedules into a single schedule group that is used by a single security policy.

Schedule groups can contain both recurring and on-time schedules. Schedule groups cannot contain other schedule groups.

Schedule expiry

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the `config firewall command` enter the command:

```
set schedule-timeout enable
```

By default, this is set to disable.

UTM profiles

Where security policies provide the instructions to the FortiGate unit as to what traffic is allowed through the device, the Unified Threat Management (UTM) profiles provide the screening that filters the content coming and going on the network. The UTM profiles enable you to instruct the FortiGate unit what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A UTM profile is a group of options and filters that you can apply to one or more firewall policies. UTM profiles can be used by more than one security policy. You can configure sets of UTM profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same UTM profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate protection profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

UTM profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure UTM profiles in the *UTM* menu and applied when creating a security policy by selecting the UTM profile type.

For more information about configuring profiles that will be used in a security policy, see the UTM chapter of the *FortiOS Handbook*.

How to use UTM profiles to monitor and protect your network

In this example, UTM profiles help you in monitoring and protecting your network from viruses, email filtering and web filtering. This example uses the default UTM profiles.

1 **Locate the security policy that allows access to the Internet (internal -> wan 1) in *Policy* > *Policy* > *Policy*.**

2 **On the Edit Policy page, select UTM and then select these options: *Enable Antivirus*, *Enable Web Filter* and *Enable Email Filter*.**

The FortiGate unit will apply the antivirus, web filter, and email filter settings to the packet if a match is found.

3 **Select *OK*.**

When packets enter the FortiGate unit's internal interface, if a packet matches the internal -> wan 1 policy, the FortiGate unit now scans for viruses and applies any web filtering and email filtering rules if there are matches as well.

4 **Go to the eicar.org web site and download the eicar test file.**

By downloading the eicar test file, you can determine that the antivirus profile is working properly, as well as to see this activity on the AV Monitor page. When attempting to download the file, a web page appears, stating that you are not permitted to download the file. This indicates that the antivirus profile is working properly.

5 **Go to *UTM Profiles* > *Monitor* > *AV Monitor* to view the virus activity that just occurred.**

On the page, you should see that the eicar test file was detected by the FortiGate unit; you can select the bar in the chart to see more details. This takes you directly to the FortiGuard Virus Encyclopedia.

6 **Go to *UTM Profiles* > *Monitor* > *Web Monitor* to view the Internet activity that is occurring on your network.**

On the page, you will see a pie chart that displays all HTTP requests and a bar chart that displays all blocked HTTP requests. If you want to view more detailed information about the blocked requests, hover your mouse over a bar; a tool-tip appears stating how many blocked requests occurred for that item. For example, for Virus, it is one blocked request because you tried to download the eicar test file.

7 **Go to *UTM Profiles* > *Monitor* > *Email Monitor* to view the email activity that is occurring on your network.**

On the page, you will see both a pie chart and a bar chart, similar to the Web Monitor page. The pie chart displays all the email activity and the bar chart displays all the blocked emails for SMTP, POP3, IMAP, and NNTP.

Security policies

Security policies control all traffic attempting to pass through the FortiGate unit, between FortiGate interfaces, zones, and VLAN subinterfaces.

This section explains what security policies are and how they affect all traffic to and from your network. This section also describes how to configure basic policies which are used as a building block to more complex policies, but they enable you to get the FortiGate unit running on the network quickly.

The following topics are included in this section:

- Security policy overview
- Policy order
- Security policies
- Creating a basic security policy

Security policy overview

Security policies are instructions the FortiGate unit uses to decide connection acceptance and packet processing for traffic attempting to pass through. When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number), and attempts to locate a security policy matching the packet.

Security policies can contain many instructions for the FortiGate unit to follow when it receives matching packets. Some instructions are required, such as whether to drop or accept and process the packets, while other instructions, such as logging and authentication, are optional. The FortiGate unit requires one security policy per traffic flow. For example, network traffic must flow from the internal network to the Internet; a security policy is created (internal interface -> external interface) that allows packets to flow freely from the Internet to the internal network, and from the internal network to the Internet.

Policy instructions may include network address translation (NAT), or port address translation (PAT), or by using virtual IPs or IP pools to translate source and destination IP addresses and port numbers.

Policy instructions may also include UTM profiles, which can specify application-layer inspection and other protocol-specific protection and logging, as well as IPS inspection at the transport layer.

You configure security policies to define which sessions will match the policy and what actions the FortiGate unit will perform with packets from matching sessions.

Sessions are matched to a security policy by considering these features of both the packet and policy:

- Source Interface/Zone
- Source Address
- Destination Interface/Zone
- Destination Address

- Schedule and time of the session's initiation
- Service and the packet's port numbers.

If the initial packet matches the security policy, the FortiGate unit performs the configured Action and any other configured options on all packets in the session.

Packet handling actions can be *ACCEPT*, *DENY*, *IPSEC* or *SSL-VPN*.

- ACCEPT policy actions permit communication sessions, and may optionally include other packet processing instructions, such as requiring authentication to use the policy, or specifying one or more UTM profiles to apply features such as virus scanning to packets in the session. An ACCEPT policy can also apply interface-mode IPsec VPN traffic if either the selected source or destination interface is an IPsec virtual interface.
- DENY policy actions block communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped, therefore it is not required to configure a DENY security policy in the last position to block the unauthorized traffic. A DENY security policy is needed when it is required to log the denied traffic, also called "violation traffic".
- IPSEC and SSL-VPN policy actions apply a tunnel mode IPsec VPN or SSL VPN tunnel, respectively, and may optionally apply NAT and allow traffic for one or both directions. If permitted by the firewall encryption policy, a tunnel may be initiated automatically whenever a packet matching the policy arrives on the specified network interface, destined for the local private network.

You need to create security policies based on how the network traffic is going to be flowing through the FortiGate unit. For example, a policy for POP3, where the email server is outside of the internal network, traffic should be from an internal interface to an external interface rather than the other way around. It is typically the user on the network requesting email content from the email server and thus the originator of the open connection is on the internal port, not the external one of the email server. This is also important to remember when view log messages as to where the source and destination of the packets can seem backwards.

If you make any changes to existing policies, those changes take effect immediately.

Security policy list details

The security policy table includes, by default, a number of columns to display information about the policy, for example, source, destination, service, and so on. You can add a number of additional columns to the table to view more information about the policies and what is in their configuration. By going to *Policy* > *Policy* > *Policy* and selecting the *Column Settings* link, you can add or remove a number of different columns of information to the policy list, and arrange their placement within the table.

Figure 5: Security policy column selection

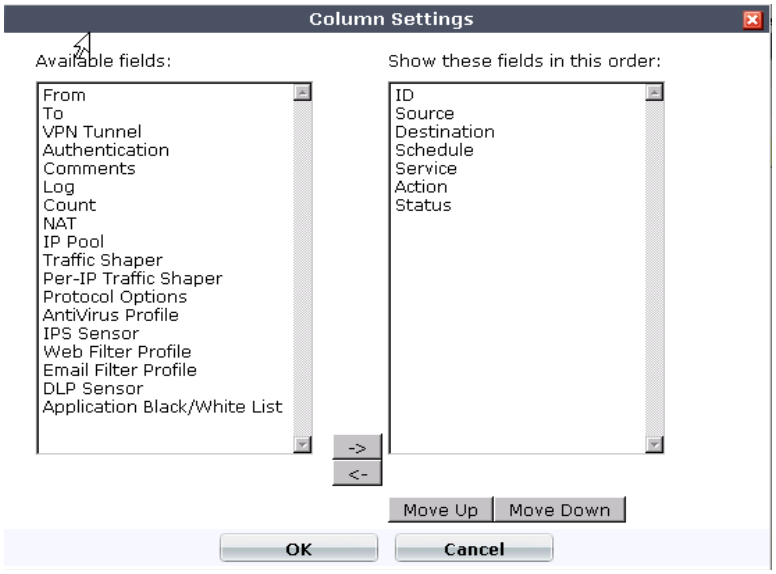

Viewing security policies

When viewing security policies in the security policy list, you can view them in either *Section View* or *Global View*. In *Section View*, policies are grouped by how the traffic is directed by interface, for example, internal -> wan1. In *Global View*, policies are listed in one large list with no groupings, referred to as interface pairings.

The FortiGate unit will automatically change the view on the policy list page to *Global View* whenever a policy containing *any* in the *Source interface/zone* or *Destination interface/zone* is created. This occurs because the FortiGate unit understands that this particular policy allows or denies traffic on any FortiGate interface, which breaks the original policy sequence order.

Policies are ordered by fixed policies (ones that contain static interfaces) with each interface pairing (for example, port1 -> port2) and each pairing has their own specific policy order, which does not cause any conflict. However, this interface pairing creates a conflict when a policy containing an ANY interface is created, because the FortiGate unit is now unable to determine which policy set to use and which, in the pair's ordering, should traffic be blocked. The FortiGate unit uses Global View to represent its own understanding of the global policy that was created, using this to help determine the action to take.

Policy order

Each time a FortiGate unit receives a connection attempting to pass through one of its interfaces, the unit searches its security policy list for a matching security policy.

The search begins at the top of the policy list and progresses in order towards the bottom. The FortiGate unit evaluates each policy in the security policy list for a match until a match is found. When the FortiGate unit finds the first matching policy, it applies the matching policy's specified actions to the packet, and disregards subsequent security policies. Matching security policies are determined by comparing the security policy and the packet's:

- source and destination interfaces
- source and destination firewall addresses

- services
- time/schedule.

If no policy matches, the connection is dropped.

As a general rule, you should order the security policy list from most specific to most general because of the order in which policies are evaluated for a match, and because only the ***first*** matching security policy is applied to a connection. Subsequent possible matches are not considered or applied. Ordering policies from most specific to most general prevents policies that match a wide range of traffic from superseding and effectively masking policies that match exceptions.

For example, you might have a general policy that allows all connections from the internal network to the Internet, but want to make an exception that blocks FTP. In this case, you would add a policy that denies FTP connections above the general policy.

Figure 6: Example: Blocking FTP — Correct policy order

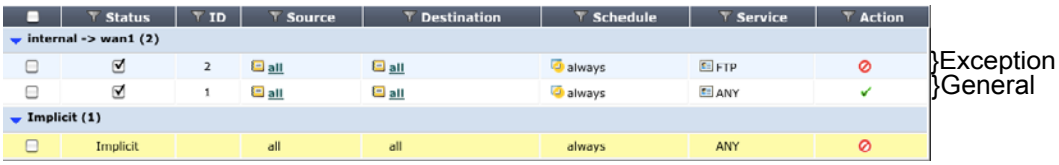

FTP connections would immediately match the deny policy, blocking the connection. Other kinds of services do not match the FTP policy, and so policy evaluation would continue until reaching the matching general policy. This policy order has the intended effect. But if you reversed the order of the two policies, positioning the general policy before the policy to block FTP, all connections, including FTP, would immediately match the general policy, and the policy to block FTP would never be applied. This policy order would not have the intended effect.

Figure 7: Example: Blocking FTP — Incorrect policy order

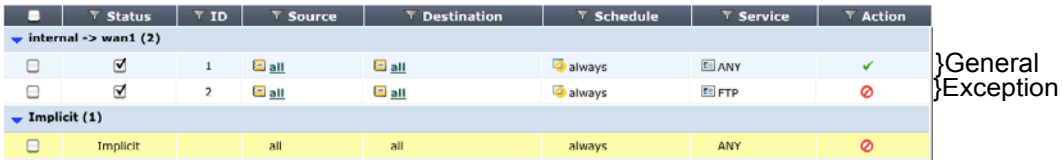

Similarly, if specific traffic requires authentication, IPsec VPN, or SSL VPN, you would position those policies above other potential matches in the policy list. Otherwise, the other matching policies would always take precedence, and the required authentication, IPsec VPN, or SSL VPN might never occur.

A default security policy may exist, which accepts all connections. You can move, disable or delete it. If you move the default policy to the bottom of the security policy list and no other policy matches the packet, the connection will be accepted. If you disable or delete the default policy and no other policy matches the packet, the connection will be dropped.

You can arrange the security policy list to influence the order in which policies are evaluated for matches with incoming traffic. When more than one policy has been defined for the same interface pair, the first matching security policy will be applied to the traffic session.

How to arrange policies

In this example, there are four policies that the FortiGate unit must use when packets enter the FortiGate unit's interface. These policies are IPsec VPN, DENY, Internet access, and an identity-based policy. You need to make sure that the policies are arranged so that the policies that are important do not get left out.

1 **On the policy list, select *Global* view to view all policies in the list.**

   By viewing the list using Global view, you can easily see all policies in the list regardless of the sections that they are in. This helps you to see where in the list you need to move the policies, without having to expand each section to view the policies.

2 **Move the IPsec VPN policy to the first line in the table.**

   You want the IPsec VPN policy to come first so that the process matches this policy first. If the IPsec VPN policy is not first, other policies would always take precedence and the authentication required for IPsec may never occur.

3 **Move the DENY policy to the third line of the table.**

   This DENY policy contains information that denies all FTP traffic.

4 **Move the identity-based policy to the fourth line in the table.**

5 **Move the Internet access policy after the identity-based policy.**

Security policies

There are many different security policies that you can configure for the FortiGate firewall. These policies include SSL VPN, wireless, and identity-based policies. With different configurations come different security policies, and each contain different information for processing the packets coming into the FortiGate unit.

The following explain each type of security policy that can be configured and the reason for configuring such a security policy.

This topic contains the following:

- Identity-based policies
- SSL VPN policies
- IPsec policies
- Accept policies
- Deny policies
- IPv6 policies
- Security policy 0
- Local-in policies

If you make any changes to existing policies, those changes take effect immediately.

Identity-based policies

If you enable *Enable Identity Based Policy* in a security policy, network users must send traffic involving a supported firewall authentication protocol to trigger the firewall authentication challenge, and successfully authenticate, before the FortiGate unit will allow any other traffic matching the security policy.

User authentication can occur through any of the following supported protocols:

- HTTP
- HTTPS
- FTP
- Telnet

Authentication can also occur through automatic login using NTLM and FSAE, to bypass user intervention.

The authentication style depends on which of these supported protocols you have included in the selected firewall services group and which of those enabled protocols the network user applies to trigger the authentication challenge. The authentication style will be one of two types. For certificate-based (HTTPS or HTTP redirected to HTTPS only) authentication, you must install customized certificates on the FortiGate unit and on the browsers of network users, which the FortiGate unit matches. For user name and password-based (HTTP, FTP, and Telnet) authentication, the FortiGate unit prompts network users to input their firewall user name and password.

For example, if you want to require HTTPS certificate-based authentication before allowing SMTP and POP3 traffic, you must select a firewall service (in the security policy) that includes SMTP, POP3 and HTTPS services. Prior to using either POP3 or SMTP, the network user would send traffic using the HTTPS service, which the FortiGate unit would use to verify the network user's certificate; upon successful certificate-based authentication, the network user would then be able to access his or her email.

In most cases, you should ensure that users can use DNS through the FortiGate unit without authentication. If DNS is not available, users will not be able to use a domain name when using a supported authentication protocol to trigger the FortiGate unit's authentication challenge.

If you do not install certificates on the network user's web browser, then network users may see an SSL certificate warning message and have to manually accept the default FortiGate certificate, which the network user's web browser may then deem as invalid.

When you use certificate authentication, if you do not specify any certificate when you create a security policy, the FortiGate unit will use the default certificate from the global settings. If you specify a certificate, the per-policy setting will override the global setting.

Authentication requires that *Action* is ACCEPT or SSL-VPN, and that you first create users, assign them to a firewall user group, and assign UTM profiles to that user group.

For additional information about identity-based-policy positioning and identity-based sub-policies, see "Identity-based security policies" on page 85.

Identity-based policy example

With this basic identity-based policy example, the security policy will allow HTTPS traffic passing from the external interface (WAN1) to the internal interface (Internal) at all times, as soon as the network user enters their user name and password. For simplicity, the policy will request the firewall authentication. This authentication can be set up for users by going to *User* > *User* > *User* and their groupings by going to *User* > *User Group* > *User Group*. For this example, the group “accounting” is used. When a user attempts to browse to a secure site, they will be prompted for their log in credentials.

To create a identity-based policy - web-based manager

1. Go to *Policy* > *Policy* > *Policy* and select *Create New*.
2. Enter the following:
3. Select *Enable Identity Based Policy*.
4. *Firewall* authentication is enabled by default.
5. Select *Add*.
6. From the *Available User Groups* list, select the *Accounting* user group and select the right arrow to move it to the *Selected User Groups* area.
7. From the *Available Services* list, select the *HTTPS* and select the right arrow to move it to the *Selected Services* area.
8. For the *Schedule*, select *Always*.
9. Select *OK*.

To create a identity-based policy - CLI

```
config firewall policy
  edit 1
    set srcintf internal
    set srcaddr 10.13.20.22
    set dstintf wan1
    set dstaddr 172.20.120.141
    set action accept
    set schedule always
    set identity-based enable
      config identity-based-policy
        edit 1
          set group accounting
          set service HTTPS
          set schedule always
        end
  end
```

SSL VPN policies

SSL VPN security policies are created for permitting SSL VPN clients, web-mode or tunnel-mode, access to the protected network behind the FortiGate unit. These security policies also contain authentication information that will authenticate the users and user group or groups.

IPsec policies

IPsec policies allow IPsec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.

For a route-based (interface mode) VPN, you do not configure an IPsec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPsec virtual interface as the source or destination interface, as appropriate.

Accept policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPsec VPN.

For information about how to configure accept policies, see “Security policy list details” on page 44.

Deny policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.

For information about how to configure DENY policies, see “Security policy list details” on page 44.

How to allow DNS queries to only one DNS server

In this example, a specific DNS server is used for all DNS queries. All other requests for DNS is not allowed. A deny security policy is used to restrict this access.

1 **In *Firewall Objects* > *Address* > *Address*, create an IP address for the DNS server.**

This address will be used for the policy that allows DNS requests from this DNS server.

2 **Create a new security policy that blocks all DNS sessions to the Internet.**

This policy would have the Action set to DENY and the Service set to DNS. In this policy, the FortiGate unit restricts all requests for any DNS queries.

3 **Create a new policy that allows access to only the DNS server.**

This policy is used by the FortiGate unit to allow DNS requests to the DNS server that is specified.

4 **Move the policies so that they are in the correct order.**

If the policies are not in the correct order, the FortiGate unit will not process the instructions properly and the policies will not work properly. The allowed policy needs to be first and the deny policy needs to come right after.

5 **Test the policies.**

You can test the policies by using diagnose debug command in the CLI or view the packet count in the Count columns of the policies. For more information about how to test and/or verify if traffic is hitting a policy, see "How to create a basic security policy for Internet access" on page 53.

IPv6 policies

IPv6 security policies are created both for an IPv6 network, and a transitional network. A transitional network is a network that is transitioning over to IPv6, but must still have access to the Internet or must connect over an IPv4 network.

These policies allow for this specific type of traffic to travel between the IPv6 and IPv4 networks. The IPv6 options for creating these policies is hidden by default. You must enable this feature in *System* > *Admin* > *Settings*.

For more information about IPv6 in FortiOS, see "Internet Protocol version 6 (IPv6)" on page 61.

Security policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPsec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPsec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate logs, you may find a log field entry indicating policyid=0. The following log message example indicates the log field policyid=0 in bold.

```
2008-10-06 00:13:49 log_id=0022013001 type=traffic
subtype=violation pri=warning vd=root SN=179089 duration=0
user=N/A group=N/A rule=0 policyid=0 proto=17 service=137/udp
app_type=N/A status=deny src=10.181.77.73 srcname=10.181.77.73
dst=10.128.1.161 dstname=10.128.1.161 src_int=N/A
dst_int="Internal" sent=0 rcvd=0 src_port=137 dst_port=137 vpn=N/A
tran_ip=0.0.0.0 tran_port=0
```

Local-in policies

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
  edit <policy_number>
    set intf <source_interface>
    set srcaddr <source_address>
    set dstaddr <destination_address>
    set action {accept | deny}
    set service <service name>
    set schedule <schedule_name>
end
```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12 using SSH on port 3 (192.168.21.77) using the Weekend schedule which defines the time the of access.

```
config firewall local-in-policy
  edit <1>
    set intf port3
    set srcaddr 192.168.21.12
    set dstaddr 192.168.21.77
    set action accept
    set service SSH
    set schedule Weekend
end
```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```
config firewall local-in-policy
  edit <policy_number>
    set status disable
end
```

Use the same commands with a `status` of `enable` to use the policy again.

Local-in policies are also supported for IPv6 by entering the command `config firewall local-in-policy6`.

Creating a basic security policy

The following describes how to configure a basic security policy as well as how to test and verify that traffic hitting the policy.

This topic includes the following:

- How to create a basic security policy for Internet access
- How to verify if traffic is hitting the basic security policy
- How to test the basic security policy

How to create a basic security policy for Internet access

The following explains how a basic security policy is created, as well as how to test and verify that the policy is working properly. Testing a policy and verifying if traffic is hitting a policy are two ways to ensure that the policy that you created is working properly.

1 **In the web-based manager, go to *Policy* > *Policy* > *Policy* and select *Create New*.**

2 **The source interface should be *internal* and the destination interface should *wan1*.**

This indicates to the FortiGate unit that the incoming packets will be coming from the internal network and proceeding to the publica network or Internet. The interfaces are also understood in reverse: packets that are coming from the outside or Internet and are destined for the internal network.

3 **The source and destination addresses should *all*.**

This is the default IP address range in *Firewall Objects* > *Addresses* > *Address*. This default IP address range indicates that any IP address is accepted within the range. This is written as 0.0.0.0/0.0.0.0.

4 **For this policy, you must choose the default *always* schedule for *Schedule*, the *ANY* service for *Service*, and the *Action* to *ACCEPT*.**

The default schedule always provides the time limitation, which is none, for the policy. A time limitation can limit the access users have to the Internet or can allow users to access resources at any time of the day or night.

5 **Select *Log Allowed Traffic* to view the traffic activity using either *Policy* > *Monitor* > *Policy Monitor*, or traffic logs. Select *OK* to save the security policy.**

You should test the policy after it has been created. To test a security policy, go to a web site; if you are able to get to the web site, the policy is working properly. You can also view the Count column on the Policy page. The Count column displays the number of packets that have recently passed through, which increases as the packets pass through the FortiGate unit.

How to test the basic security policy

After a security policy has been configured, you can test to see if the policy is working. This should be done after you create a security policy so that you can modify the policy's settings, if required, before backing up the configuration. You should always back up the configuration after making modifications to the FortiGate configuration; by doing so, you will have a current configuration whenever you need it.

1 **On a computer that is on the internal network, open a web browser and access any web site.**

You should be able to get to that web site.

2 **If you are unable to get to a web site, use the following to help troubleshoot the problem:**

- Is the policy order correct?
- Using the diag debug flow command, see if traffic is hitting the policy. If not, use the diag sniffer command to determine what is going on
- View the Count column; if no number appears, traffic is not hitting the policy.

3 **After troubleshooting the problem, browse to a web site and if you can access it, and then save the current configuration.**

How to verify if traffic is hitting the basic security policy

After configuring a security policy, you will want to verify that it is working properly. The following explains how to verify that traffic is hitting the basic security policy that you configured in "How to verify if traffic is hitting the basic security policy" on page 54.

1. **In the web-based manager, go to *Policy* > *Policy* > *Policy* and locate the internal to wan1 policy.**
2. **In the Count column, verify that there is packets hitting the security policy.**

   The Count column displays the amount of packets that are hitting the security policy. In the beginning this count will be low, be increase as the packets come through the FortiGate unit.
3. **Go to *Policy* > *Monitor* > *Policy Monitor* to view the security policy.**

   On the Policy Monitor page, you can see the active sessions, bytes or packets that are occurring from the bar chart and table. By selecting the bar within the chart, you can view more detailed information.
4. **Go to the CLI, log in, and use `diag debug flow` commands to show traffic is hitting the security policy.**

   The `diag debug flow` commands show packet flow through the FortiGate unit. The following is an example of what the information gives when you use the `diag debug flow` commands to see if traffic is hitting a policy.

```
diagnose debug enable
diagnose debug flow show console enable
diagnose debug flow filter add 192.168.1.110
diagnose debug flow trace start 50

id=36871 trace_id=1 msg="vd-root received a packet(proto=6,
   192.168.1.110:3152->172.16.100.148:80) from internal."
id=36871 trace_id=1 msg="allocate a new session-0000724b"
id=36871 trace_id=1 msg="find a route: gw-172.20.120.2 via wan1"
id=36871 trace_id=1 msg="find SNAT: IP-172.20.120.11, port-
   40156"
id=36871 trace_id=1 msg="Allowed by Policy-3: SNAT"
id=36871 trace_id=1 msg="SNAT 192.168.1.110-
   >172.20.120.11:40156"
```

Monitoring firewall traffic

You can easily monitor the network traffic on the FortiGate firewall from either the Dashboard or the Monitor menus in the Policy and Firewall Objects menus. By using these monitors, you can understand how to improve your firewall, or resolve issues.

The following explains the various features that you can use to monitor firewall traffic.

The following topics are included in this section:

- Session tables
- Monitoring security policy traffic activity

Session tables

Firewall session tables include entries to record source and destination IP addresses and port numbers. For each packet received by a FortiGate unit, it references the session table for a match. Packets of an established session are checked against the session table continually throughout the communication. The performance of depends on the performance of processing session table.

Firewall sessions clear from the table based on the timeout, that is, Time-to-live (TTL) setting. Equally, a completely inactive session with no FIN or RESET will be flushed by the by the session TTL timer. Sessions are not closed based on FIN or a RESET. A FIN that is acknowledged with a FIN ACK would slush the session.

Viewing session tables in the web-based manager

Firewall sessions are viewable in the web-based manager using the *Top Sessions* widget. If this widget is not on the Dashboard, select the *Widget* link at the top of the web-based manager and select it from the pop-up dialog box.

While this view shows a graph of the connecting users and IP addresses, double-clicking on a bar in the graph will display the complete session information for that user.

You can clear a session from the table by scrolling to the right and selecting the delete icon for a given session.

Sessions Monitor

Session information display in *Policy* > *Monitor* > *Session Monitor*. You can delete sessions, refresh so that you are viewing current sessions, and you can also filter the session information on the page as well. Filtering allows you to view specific information. For example, you want to view only TCP sessions.

Session Monitor page

Displays the sessions that are currently being monitored by the unit.

Refresh	Select to refresh the information in the list.

Filter Settings	Select to filter the information on the page. *Filters* appears automatically after selecting *Filter Settings*, below the column headings. Use to configure filter settings. **Note:** *Filter Settings* configures all filter settings. Filter icons are used to configure filter settings within that column. To apply a filter setting, select the plus sign beside *Add new filter* and then select and enter the information required. Repeat to add other filter settings. To modify the settings, select *Change* beside the setting and edit the settings. To clear all filters settings. Select the icon beside *Clear all filters*. To use a filter icon to filter settings within a column, select the filter icon in the column. *Filters* appears. Within *Filters*, configure the settings for that column.
IPv4	Select to display only IPv4 addresses.
IPv6	Select to display only IPv6 addresses.
Both	Select to display both IPv4 and IPv6 addresses.
Total Concurrent Sessions: <number>/ New Sessions per Second: <number>	Indicates the total number of concurrent sessions, as well as new sessions that are occurring each second.
Page Controls	Use to navigate through the list.
Total: <number>	The total number of current sessions.
#	The number of the session within the list.
Protocol	The service protocol of the connection, for example, UDP.
Src Address	The source IP address of the connection.
Src Port	The source port of the connection.
Src NAT IP	The source NAT IP address.
Src NAT Port	The source NAT IP port.
Dst Address	The destination address of the connection.
Dst Port	The destination port of the connection.
Policy ID	The security policy identification number.
Expiry (sec)	The time, in seconds, before the connection expires.
Duration (sec)	The duration, in seconds, of the session.
Delete	Select to remove a session from within the list.

Viewing session tables in the CLI

Session tables and information is also viewable from the CLI. More information on sessions are available from the CLI where various diagnose commands reveal more granular data. To view the session information enter the following CLI command:

```
diagnose sys session list
```

Output will look something similar to:

```
session info: proto=17 proto_state=01 duration=121 expire=58
   timeout=0 flags=00000000 sockflag=00000000 sockport=0
   av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 hakey=0
policy_dir=0 tunnel=/
state=may_dirty br
statistic(bytes/packets/allow_err): org=63/1/1 reply=133/1/1
   tuples=2
orgin->sink: org pre->post, reply pre->post dev=6->2/2->6
   gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 172.20.120.85:51167-
   >8.8.8.8:53(0.0.0.0:0)
hook=post dir=reply act=noop 8.8.8.8:53-
   >172.20.120.85:51167(0.0.0.0:0)
misc=0 policy_id=3 id_policy_id=0 auth_info=0 chk_client_info=0
   vd=0
serial=000171db tos=ff/ff app_list=0 app=0
dd_type=0 dd_rule_id=0
per_ip_bandwidth meter: addr=172.20.120.85, bps=1984
total session 189
```

To clear a session enter the following command:

```
diagnose sys session clear
```

State	Meaning
log	Session is being logged
local	Session is originating from, or destined for, a local stack.
ext	Session is created by a firewall session helper.
may_dirty	Session is created by a policy. For example, the session for FTP channel control will have this state but the FTP data channel will not.
ndr	Session will be checked by an IPS signature.
nds	Session will be checked by an IPS anomaly.
br	Session is being bridged, that is, in transparent mode.
npu	Session will possibly be offloaded to NPU.
wccp	Session is handled by WCCP.

Proto_state fields: TCP

The proto_state field value has two digits. This is because the FortiGate unit keeps track of the original direction and the reply direction.

State	Value	Expire Timer Default (seconds)
NONE	0	10
ESTABLISHED	1	3600
SYN_SENT	2	120

SYN & SYN/ACK	3	60
FIN_WAIT	4	120
TIME_WAIT	5	120
CLOSE	6	10
CLOSE_WAIT	7	120
LAST_ACK	8	30
LISTEN	9	120

Proto_state fields: SCTP

State	Value	Expire Timer Default (seconds)
SCTP_S_NONE	0	60
SCTP_S_ESTABLISHED	1	3600
SCTP_S_CLOSED	2	10
SCTP_S_COOKIE_WAIT	3	5
SCTP_S_COOKIE_ECHOED	4	10
SCTP_S_SHUTDOWN_SENT	5	30
SCTP_S_SHUTDOWN_RECD	6	30
SCTP_S_ACK_SENT	7	3
SCTP_S_MAX	8	120

Proto_state fields: UDP

UDP is a sessionless protocol, however the FortiGate unit still monitors two different states:

- Reply Not Seen - 0
- Reply Seen - 1

In the example output below, a state of 00, the UDP packet has been seen and a session will be created, but no reply packet has been seen:

```
session info: proto=17 proto_state=00 expire=179 timeout=3600
   use=3
```

In this example, the UDP packet has been seen and a session created. Reply packets have also been seen:

```
session info: proto=17 proto_state=01 expire=22 timeout=3600
   use=3
```

Proto_state field for ICMP

There are no states for ICMP traffic; it will always appear as `proto_state=00`.

Monitoring security policy traffic activity

The Policy Monitor page provides information about the activity of security policies. This activity can be viewed at a high level, or in much more detail, by drilling down to get more specific information.

The Policy Monitor page allows you to view the information in either a graphical format, or in a table. The graphical format, or chart, provides an easy and user-friendly view of the traffic activity that is occurring. The chart also provides a way to drill-down to more information; you can view this information by selecting on a bar within the chart. The drill-down information can be displayed by source address or destination address or by destination port.

Below the chart, a table provides information as well about each policy include the type of action the policy

Policy Monitor page

Displays information about the security policy traffic occurring on the unit.

Tip: View additional and more detailed information by selecting a bar within the chart.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Policy Usage	Displays the top security policy usage in a bar chart.
Report By	Select to view information by the current active sessions, bytes or packets.

(Table explaining detailed information about the top policy usage)

Policy ID	The security policy identification number.
Source Interface/Zone	The source address or zone used within that security policy.
Destination Interface/Zone	The destination address of zone used within that security policy.
Action	The type of action that is specified in the security policy. For example *Action* is set to *DENY*. The action displays as an icon; for example, a green check mark is *ALLOW*.
Bytes	The number of bytes used by the security policy. This is reflected in the bar chart.
Packets	The number of packets.

Internet Protocol version 6 (IPv6)

This section explains IPv6 in FortiOS. This section does not explain IPv6 in its entirety, only a high-level summary of IPv6 and how IPv6 is supported in FortiOS. For any additional information about IPv6, see the ipv6.com web site.

The following topics are included in this section:

- What is IPv6?
- IPv6 in FortiOS
- Dual stack routing configuration
- IPv4 tunneling configuration
- Remotely connecting to an IPv6 network over the Internet

What is IPv6?

Internet Protocol version 6 (IPv6) is the next-generation version of IP addressing. This updated version of IP addressing provides many advances, such as more routing efficiency and eliminating the need for NAT. IPv6 also provides better security and mobility support, as well as stateless auto-reconfiguration of hosts which allows IPv6 hosts to automatically configure when connected to a routed IPv6 network.

IPv6 uses 128-bit addressing, which is written in hexadecimal digits separated by a colon. For example, 2001:DB8::6334. This revised version of IP addressing has the potential to provide trillions and trillions of addresses, or an address for each device on the Internet.

For IPv6 address examples, documents use the IPv6 special address 2001:DB8::/32 to indicate that the address is an example. This is stated in RFC 3849. For more information about the specific addresses that are used in IPv6, see ipv6.com.

IPv6 in FortiOS

By default, the FortiGate unit is not enabled to use IPv6 options and settings; however, they are there. To enable IPv6, go to *System* > *Admin* > *Settings* and select *IPv6 Support on GUI*. When enabled, you can use IPv6 addressing on any of the address-dependant components of the FortiGate unit, including security policies, interface addressing and DNS servers. IPv6 addressing can be configured on the web-based manager and in the CLI.

There are many different features that FortiOS supports in IPv6. The following is what FortiOS supports in IPv6:

- Static routing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- DNS
- Network interface addressing
- Packet and network sniffing
- IPsec VPN
- SSL VPN
- UTM protection

- Routing access lists and prefix lists
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Security policies
- Authentication
- IPv6 over SCTP
- UTM protection
- NAT/Route and Transparent mode
- Logging and reporting
- SNMP
- Virtual IPs and groups
- IPv6-specific troubleshooting, such as ping6

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network.

IPv6 works almost the same as IPv4 in FortiOS. The only main difference is the IP addresses, since you are using IPv6 addressing instead of IPv4. There is also no NAT, unless you are configuring a dual stack routing or IPv4 tunnelling configuration.

Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

Dual stack routing configuration

A dual stack routing configuration implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. The FortiOS dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary.

Devices that are on this type of configured network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses. The dual stack architecture of FortiOS provides all the features that you need for protecting your network, such as UTM security for the traffic, and routing.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunneling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

IPv4 tunneling configuration

In an IPv4 tunneling configuration, IPv6 packets are encapsulated within IPv4 headers, which carry these IPv6 packets over IPv4 tunnels. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

Remotely connecting to an IPv6 network over the Internet

Similar to the IPv4 tunneling configuration, FortiOS supports IPv6 tunneling over IPv4 across the Internet between two IPv6 networks that are protected by FortiGate units.

All traffic between the IPv6 networks are tunnelled over IPv4, which in this case is the Internet. Each FortiGate unit extracts the IPv6 traffic from the IPv4 tunnel and traffic on the internal networks uses IPv6.

In FortiOS, you configure this type of network configuration using IPsec VPN because IPv6 is supported for IPsec VPNs. The VPN provides higher security for the data transmitted between the IPv6 networks. This configuration includes an interface-based IPsec VPN between IPv6 interfaces on each FortiGate unit.

Advanced FortiGate firewall concepts

The FortiGate firewall has advanced firewall component options, which allows for greater flexibility when these advanced options are needed to help with your growing network. These advanced firewall components include traffic shaping, QoS and identity-based policies.

The following topics are included in this section:

- Central NAT table
- Stateful inspection of SCTP traffic
- Port pairing
- Blocking port 25 to email server traffic
- Blocking HTTP access by IP
- ICMP packet processing
- Adding NAT security policies in Transparent mode
- Adding a static NAT virtual IP for a single IP address and port
- Double NAT: combining IP pool with virtual IP
- Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping
- Traffic shaping and per-IP traffic shaping
- Endpoint Security
- Logging traffic
- Quality of Service (QoS)
- Identity-based security policies

Central NAT table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The NAT table also functions in the same way as the security policy table. That is, the FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. The NAT policies can be rearranged within the policy list as well, the same way as security policies. NAT policies are applied to network traffic after a security policy.

To view the Central NAT configuration page, and use them in a security policy, you need to first enable it.

To enable Central NAT - web-based manager

1 Go to *System* > *Admin* > *Settings*.

2 In the Display Options on GUI section, select the check box beside *Central NAT table*.

3 Select *Apply*.

To enable Central NAT - CLI

```
config system global
  set gui-central-nat-table
end
```

NAT policies are created in the web-based manager by going to *Policy* > *Policy* > *Central NAT Table*. The NAT policies are enabled when you configure the security policy by selecting the *Use Central NAT Table* option.

NAT policies are created in the CLI by using the commands under `config firewall central-nat`. To enable the policies use the commands

```
config security policy
  edit <policy_number>
    set central-nat enable
end
```

Central NAT Table configuration settings

To configure the Central NAT table, go to *Policy* > *Policy* > *Central NAT Table* and select *Create New*.

New NAT page	
Source Address	Select the source IP address from the drop-down list. You can optionally create a group of source IP addresses when you select *Multiple* in the drop-down list. You can also create a new source IP address when you select *Create New* in the drop-down list.
Translated Address	Select the dynamic IP pool from the drop-down list.
Original Source Port	Enter the source port that the address is originating from.
Translated Port	Enter the translated port number. The number in the *From* field must be greater than the lower port number that is entered in the *To* field.

Stateful inspection of SCTP traffic

The Stream Control Transmission Protocol (SCTP) is a transport layer protocol similar to TCP and UDP. SCTP is designed to provide reliable, in-sequence transport of messages with congestion control. SCTP is defined in RFC 4960.

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a reliable transport protocol that runs on top of a connectionless packet network (IP). SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure-detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Configuring FortiGate SCTP filtering

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to ANY. FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism

- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPsec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

Adding an SCTP custom service

This example creates a custom SCTP service that accepts SCTP traffic using destination port 2905. SCTP port number 2905 is used for SS7 Message Transfer Part 3 (MTP3) User Adaptation Layer (M3UA) over IP.

To add the SCTP custom service - web-based manager

1 Go to *Firewall Objects* > *Service* > *Custom* and select *Create New*.

2 Enter the following and select *OK*.

Name	M3UA_service
Protocol Type	TCP/UDP/SCTP
Protocol	SCTP
Source Port (Low)	1
Source Port (High)	65535
Destination Port (Low)	2905
Destination Port (High)	2905

To add the SCTP custom service - CLI

```
config firewall service custom
  edit M3UA_service
    set protocol TCP/UDP/SCTP
    set sctp-portrange 2905
  end
```

Adding an SCTP policy route

You can add policy routes that route SCTP traffic based on the SCTP source and destination port as well as other policy route criteria. The SCTP protocol number is 132.

The following example directs all SCTP traffic with SCTP destination port number 2905 to the next hop gateway at IP address 1.1.1.1.

To add the policy route - web-based manager

1 Go to *Router* > *Static* > *Policy Route*.

2 Select *Create New*.

3 Enter the following information and select *OK*.

Protocol	132
Incoming interface	internal
Source address / mask	0.0.0.0 0.0.0.0
Destination address / mask	0.0.0.0 0.0.0.0
Destination Ports	From 2905 to 2905

Force traffic to:

Outgoing interface	external
Gateway Address	1.1.1.1

To add the policy route - CLI

```
config router policy
  edit 1
    set input-device internal
    set src 0.0.0.0 0.0.0.0
    set dst 0.0.0.0 0.0.0.0
    set output-device external
    set gateway 1.1.1.1
    set protocol 132
    set start-port 2905
    set end-port 2905
  end
```

Changing the session time to live for SCTP traffic

Use the following command to change the session timeout for SCTP protocol M3UA on port 2905 to 3600 seconds.

```
config system session-ttl
  config port
  edit 1
    set protocol 132
    set start-port 2905
    set end-port 2905
    set timeout 3600
  end
end
```

Port pairing

Port pairing is an option in Transparent mode to bind two ports together. In doing this, you can create security policies that regulate traffic only between two specific ports, VLANs or VDOMs. In its simplest form, this enables an administrator to create security policies that are only between these two ports. Traffic is captured between these ports. No other traffic can enter DNS services or leave a port pairing.

For example, a FortiGate unit has three ports, where port 1 and port 2 are paired together, because the two networks only need to communicate with each other. If packet arrives on port 1, the FortiGate unit needs to figure out whether the packet goes to port 2 or port 3. With port pairing configured, it is more simple. If packet arrives on port 1, then the FortiGate automatically directs the packet to port 2. The opposite is also true in the other direction. This can be ideal when to groups only need to transfer data between each other.

Figure 8: Port pairing

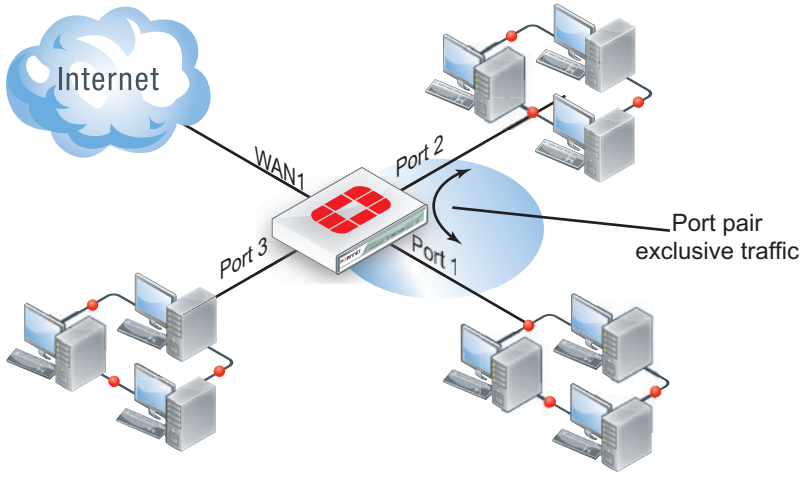

To configure port pairing - web-based manager

1 Go to *System* > *Network* > *Interface*.
2 Select the arrow beside *Create New*, and select *Port Pair*.
3 Enter a *Name* for the port pair.
4 Select the physical or virtual ports from the *Available Members* list and select the right-facing arrow to add the ports to the *Selected Members* list.

   There can be only two ports added.
5 Select *OK*.

To configure port pairing - CLI

```
config system port-pair
  edit <pair_name>
    set member <port_names>
end
```

When configuring security policies with the port pairs, selecting the *Source Interface* automatically populates the *Destination Interface*, and vice versa. All other aspects of the security policy configuration remains the same.

Blocking port 25 to email server traffic

Port 25 is the default port for SMTP traffic. Certain types of malware can install themselves on an unsuspecting user's computer and send spam using its own email server. By blocking port 25, this prevents a host system, and potentially your network or company, from being deemed a spam source.

This does, however limit your corporation from using a web server. You have a few options for this:

- if the email server is on a dedicated port, such as a DMZ port, security policies can ensure no traffic goes out from this port except the email server.
- Block all traffic on port 25 except the specific address of the email server.

Dedicated traffic

This example show the steps to ensure only traffic exits from the DMZ where the email server is connected. The internal port is connected to the internal network and the WAN1 port connects to the Internet.

First, create a security policy that will not allow any traffic through port 25 from the internal interface, which connects to the internal network. Place this policy at the top of the security policy list.

To block traffic on port 25 - web-based manager

1 Go to *Policy* > *Policy* > *Policy* and select *Create New*.

2 Set the following options and select *OK*.

Source Interface	Internal
Source Address	ALL
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	DENY
Comments	Prevent Malware spam.

You may also want to enable *Log Violation Traffic* to see if there is any potential malware or other user sending email using the non-corporate email server.

To block traffic on port 25 - CLI

```
config security policy
  edit <policy_number>
    set srcintf Internal
    set srcaddr all
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action deny
    set comment "Prevent Malware spam."
end
```

Next, create a security policy for the email server, IP address 10.10.11.29 that only allows SMTP traffic from the email server on port 25.

To allow traffic on port 25 for the email server - web-based manager

3 Go to *Policy* > *Policy* > *Policy* and select *Create New*.

4 Set the following options and select *OK*.

Source Interface	DMZ
Source Address	10.10.11.29
Destination Interface	WAN1
Destination Address	ALL

Schedule	ALWAYS
Service	SMTP
Action	ACCEPT

To allow traffic on port 25 for the email server- CLI

```
config security policy
  edit <policy_number>
    set srcintf dmz
    set srcaddr 10.10.11.29
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action allow
  end
```

Restricting traffic on port 25

This example shows how to limit traffic on port 25 on the wan port to only traffic from the email server. The web server's address is 10.10.10.29.

To allow traffic on port 25 for the email server - web-based manager

1 Go to *Policy* > *Policy* > *Policy* and select *Create New*.

2 Set the following options and select *OK*.

Source Interface	INTERNAL
Source Address	10.10.10.29
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	ACCEPT

To allow traffic on port 25 for the email server- CLI

```
config security policy
  edit <policy_number>
    set srcintf internal
    set srcaddr 10.10.10.29
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action allow
  end
```

Next, add a deny security policy that blocks all SMTP traffic from the Internal port to the WAN1 port. Ensure this policy is directly after the policy created above.

To block SMTP traffic on port 25 for the rest of the company - web-based manager

3 Go to *Policy* > *Policy* > *Policy* and select *Create New*.

4 Set the following options and select *OK*.

Source Interface	INTERNAL
Source Address	ALL
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	DENY

To block SMTP traffic on port 25 for the rest of the company - CLI

```
config security policy
  edit <policy_number>
    set srcintf internal
    set srcaddr all
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action deny
  end
```

Blocking HTTP access by IP

To block a web site using the IP, create a URL filter entry, using the additional information below. Note that this is only effective with HTTP or FortiGate units running Deep Inspection.

You need to create two URL filter entries. The first filter only allowing a text string containing two or more sets of text separated by a period. This is to match the various domain possibilities for web sites, for example:

- example.org
- www.example.com
- www.example.co.jp

The second filter blocks any IP address lookup.

To add the URL filter entries

1 Go to *UTM Profiles* > *Web Filter* > *URL Filter*.
2 Select *Create New* to add a filter group, give it a name and select *OK*.
3 Select *Create New* for a new filter.
4 Enter the *URL* of `^([a-z0-9-]+\.){1,}[a-z]+`
5 Set the *Type* to *Regex*.
6 Set the *Action* to *Allow*.
7 Select *OK*.
8 Select *Create New*.
9 Enter the *URL* of `[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}`

10 Set the *Type* to *Regex*.

11 Set the *Action* to *Block*.

12 Select *OK*.

Position these at the end of the URL filter list so that any exemptions or blocks before that are still effective.

Both of these filter entries are required. If you only enter the second one, the FortiGate unit will also catch a URL lookup as they both behave in a similar fashion after the URL is resolved to an IP. The first entry is needed to break out of the URL filter and allow the web site before it does the second check if they entered text.

ICMP packet processing

ICMP messages are used to relay feedback to the traffic source that the destination IP is not reachable. ICMP message types are:

- ICMP_ECHO
- ICMP_TIMESTAMP
- ICMP_INFO_REQUEST
- ICMP_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

Common error messages include:

- destination unreachable messages
- time exceeded messages
- redirect messages

For example, a security policy that allows TFTP traffic through the FortiGate unit. User1 (192.168.21.12) attempts to connect to the TFTP server (10.11.100.1), however, the UDP port 69 has not been opened on the server. The corresponding sniffer trace occurs:

```
diagnose sniffer packet any "host 10.11.100.1 or icmp 4"
  3.677808 internal in 192.168.21.12.1262 -> 10.11.100.1.69: udp 20
  3.677960 wan1 out 192.168.21.12.1262 -> 10.11.100.1.69: udp 20
  3.678465 wan1 in 10.11.100.1.132 -> 192.168.21.12: icmp: 10.11.100.1
      udp port 69 unreachable
  3.678519 internal out 10.11.100.1 -> 192.168.21.12: icmp:
      192.168.182.132 udp port 69 unreachable
```

Adding NAT security policies in Transparent mode

Similar to operating in NAT mode, when operating a FortiGate unit in Transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.

- Add IP pools as required for source address translation

For NAT firewall policies to work in NAT mode you must have two interfaces on two different networks with two different subnet addresses. Then you can create firewall policies to translate source or destination addresses for packets as they are relayed by the FortiGate unit from one interface to the other.

A FortiGate unit operating in Transparent mode normally has only one IP address, the management IP. To support NAT in Transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

In the example shown in Figure 9, all of the PCs on the internal network (subnet address 192.168.1.0/24) are configured with 192.168.1.99 as their default route. One of the management IPs of the FortiGate unit is set to 192.168.1.99. This configuration results in a typical NAT mode firewall. When a PC on the internal network attempts to connect to the Internet, the PC's default route sends packets destined for the Internet to the FortiGate unit internal interface. Similarly on the DMZ network (subnet address 10.1.1.0/24) all of the PCs have a default route of 10.1.1.99.

This example describes adding an internal to WAN1 security policy to relay these packets from the internal interface out the WAN1 interface to the Internet. Because the WAN1 interface does not have an IP address of its own, you must add an IP pool to the WAN1 interface that translates the source addresses of the outgoing packets to an IP address on the network connected to the wan1 interface.

The example describes adding an IP pool with a single IP address of 10.1.1.201. So all packets sent by a PC on the internal network that are accepted by the Internal to WAN1 policy leave the WAN1 interface with their source address translated to 10.1.1.201. These packets can now travel across the Internet to their destination. Reply packets return to the WAN1 interface because they have a destination address of 10.1.1.201. The Internal to WAN1 NAT policy translates the destination address of these return packets to the IP address of the originating PC and sends them out the internal interface to the originating PC.

Use the following steps to configure NAT in Transparent mode

- Add two management IPs
- Add an IP pool to the WAN1 interface
- Add an Internal to WAN1 security policy

You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

Figure 9: Example NAT in Transparent mode configuration

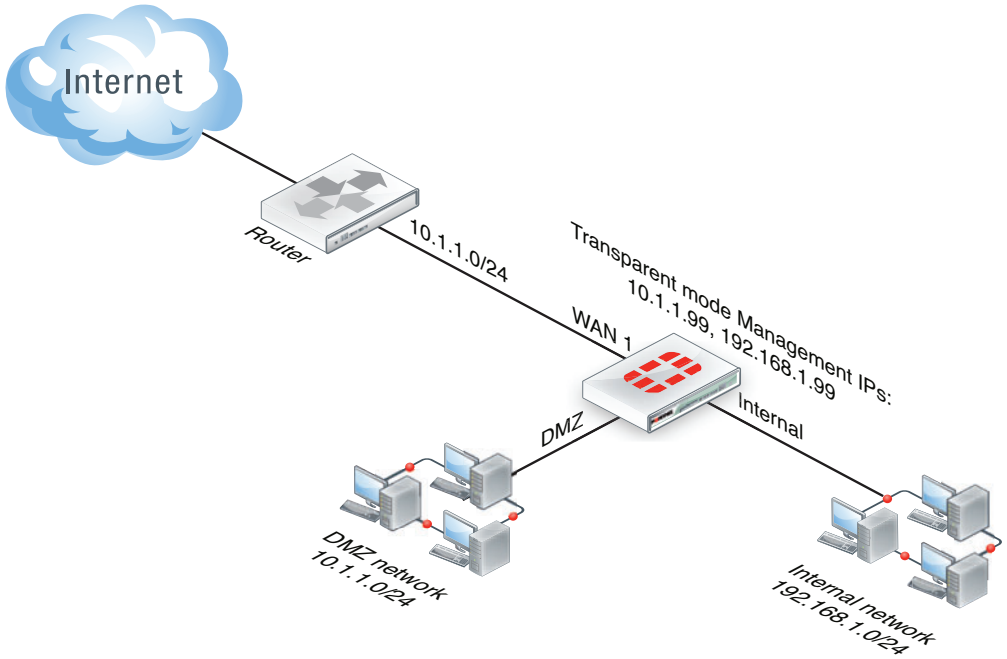

To add a source address translation NAT policy in Transparent mode

1 Enter the following command to add two management IPs.

The second management IP is the default gateway for the internal network.

```
config system settings
  set manageip 10.1.1.99/24 192.168.1.99/24
end
```

2 Enter the following command to add an IP pool to the WAN1 interface:

```
config firewall ippool
  edit nat-out
    set interface "wan1"
    set startip 10.1.1.201
    set endip 10.1.1.201
  end
```

3 Enter the following command to add an Internal to WAN1 security policy with NAT enabled that also includes an IP pool:

```
config security policy
  edit 1
    set srcintf "internal"
    set dstintf "wan1"
    set scraddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set ippool enable
    set poolname nat-out
  end
```

Adding a static NAT virtual IP for a single IP address and port

In this example, the wan1 interface of the FortiGate unit is connected to the Internet and the Internal interface is connected to the DMZ network. The IP address 192.168.37.4 on port 80 on the Internet is mapped to 10.10.10.42 on port 8000 on the private network. Attempts to communicate with 192.168.37.4 from the Internet are translated and sent to 10.10.10.42 by the FortiGate unit. The computers on the Internet are unaware of this translation and see a single computer at 192.168.37.4 rather than a FortiGate unit with a private network behind it.

Figure 10: Static NAT virtual IP for a single IP address example

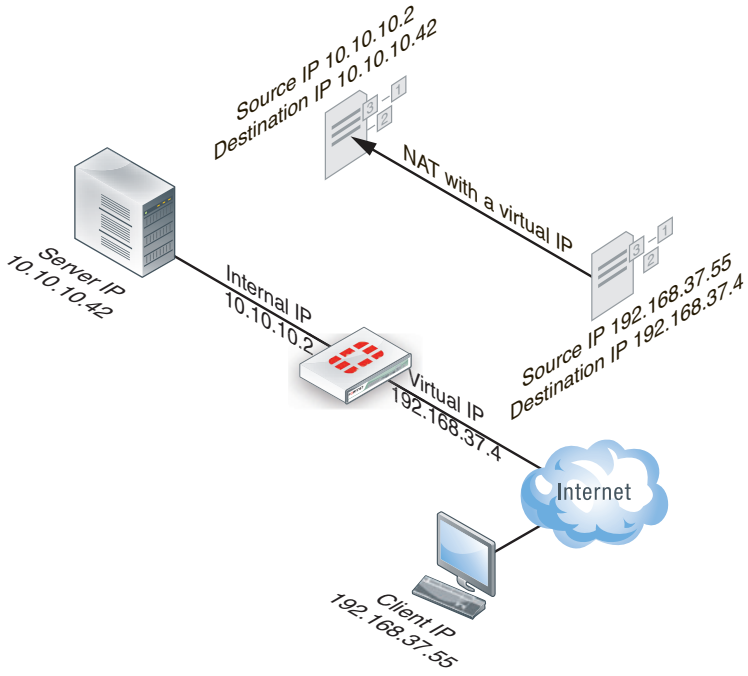

To add a static NAT virtual IP for a single IP address and port - web-based manager

1 Go to *Firewall Objects* > *Virtual IP* > *Virtual IP*.

2 Select *Create New*.

3 Complete the following and select *OK*.

Name	static_NAT
External Interface	wan1
Type	Static NAT
External IP Address/Range	192.168.37.4.
Mapped IP Address/Range	10.10.10.42
Port Forwarding	Selected
Protocol	TCP
External Service Port	80
Map to Port	8000

To add a static NAT virtual IP for a single IP address and port - CLI

```
config firewall vip
  edit static_NAT
    set extintf wan1
    set type static-nat
    set extip 192.168.37.4
    set mappedip 10.10.10.42
    set portforward enable
    set extport 80
    set mappedport 8000
  end
```

Add a external to dmz1 security policy that uses the virtual IP so that when users on the Internet attempt to connect to the web server IP address packets pass through the FortiGate unit from the external interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the external IP to the DMZ network IP address of the web server.

To add a static NAT virtual IP for a single IP address to a security policy - web-based manager

1 Go to *Policy* > *Policy* > *Policy* and select *Create New*.

2 Complete the following:

Source Interface/Zone	wan1
Source Address	All
Destination Interface/Zone	Internal
Destination Address	static_nat
Schedule	always
Service	HTTP
Action	ACCEPT

3 Select *NAT*.

4 Select *OK*.

To add a static NAT virtual IP for a single IP address to a security policy - CLI

```
config security policy
  edit 1
    set srcintf wan1
    set dstintf internal
    set srcaddr all
    set dstaddr static_nat
    set action accept
    set schedule always
    set service ANY
    set nat enable
end
```

Double NAT: combining IP pool with virtual IP

In this example, a combination of virtual IPs, IP pools and security policies will allow the local users to access the servers on the DMZ. The example uses a fixed port and IP pool to allow more than one user connection while using virtual IP to translate the destination port from 8080 to 80. The security policy uses both the IP pool and the virtual IP for double IP and/or port translation.

For this example:

- Users in the 10.1.1.0/24 subnet use port 8080 to access server 172.20.120.1.
- The server's listening port is 80.
- Fixed ports must be used.

Figure 11: Double NAT

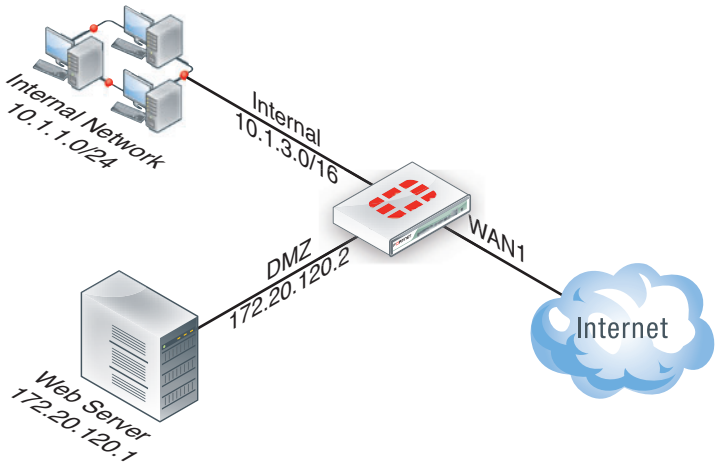

To create an IP pool - web-based manager

1 Go to *Firewall Objects* > *Virtual IP* > *IP Pool*.
2 Select *Create New*.
3 Enter the *Name* `pool-1`.
4 Enter the *IP Range/Subnet* `10.1.3.1-10.1.3.254`.
5 Select *OK*.

To create an IP pool - CLI

```
config firewall ippool
  edit pool-1
  set startip 10.1.3.1
  set endip 10.1.3.254
end
```

Next, create the virtual IP with port translation to translate the user internal IP used by the network users to the DMZ port and IP address of the server.

To create a Virtual IP with port translation - web-based manager

1 Go to *Firewall Objects* > *Virtual IP* > *Virtual IP*.
2 Select *Create New*.

3 Enter the following information and select *OK*.

Name	server-1
External Interface	Internal
Type	Static NAT
External IP Address/Range	172.20.120.1 **Note:** This address is the same as the server address.
Mapped IP Address/Range	172.20.120.1
Port Forwarding	Enable
Protocol	TCP
External Service Port	8080
Map to Port	80

To create a Virtual IP with port translation - CLI

```
config firewall vip
  edit server-1
    set extintf internal
    set type static-nat
    set extip 172.20.120.1
    set mappedip 172.20.120.1
    set portforward enable
    set extport 80
    set mappedport 8080
end
```

Add an internal to DMZ security policy that uses the virtual IP to translate the destination port number and the IP pool to translate the source addresses.

To create the security policy - web-based manager

1 Go to *Policy* > *Policy* > *Policy* and select *Create New*.

2 Complete the following and select *OK*:

Source Interface/Zone	internal
Source Address	all
Destination Interface/Zone	dmz
Destination Address	server-1
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select
Dynamic IP Pool	Select, and select the *pool-1* IP pool.

To create the security policy - CLI

```
config security policy
```

```
    edit 1
      set srcintf internal
      set dstintf dmz1
      set srcaddr all
      set dstaddr server-1
      set action accept
      set schedule always
      set service HTTP
      set nat enable
      set ippool enable
      set poolname pool-1
end
```

Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping

VIP addresses are typically used to map external (public) to internal (private) IP addresses for Destination NAT (DNAT).

This example shows how to use VIP ranges to perform source NAT (SNAT) with a static 1-to-1 mapping from internal to external IP addresses. This is similar to using an IP pool with the advantage of having predictable and static 1-to-1 address mapping.

Figure 12: Network diagram

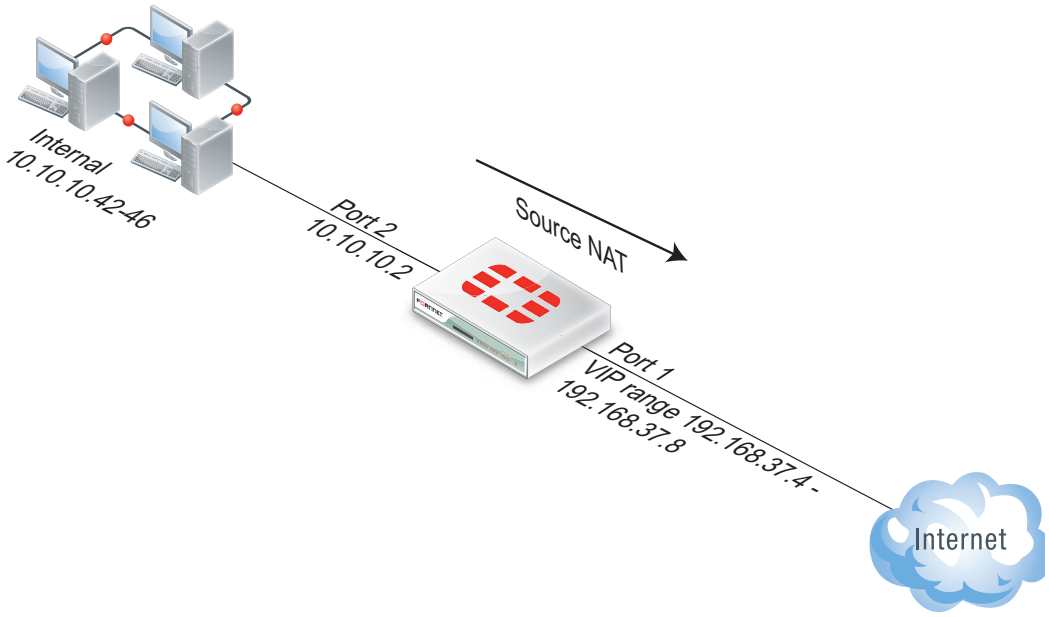

This example will associate each internal IP address to one external IP address for the Source NAT (SNAT) translation.

Using the diagram above, the translations will look like the following:

Traffic from Source IP	Translated to Source IP (SNAT)
10.10.10.42	192.168.37.4
10.10.10.43	192.168.37.5

...	...
10.10.10.46	192.168.37.8

First, configure the virtual IP.

To configure the virtual IP - web-based manager

1 Go to *Firewall Objects* > *Virtual IP* > *Virtual IP* and select *Create New*.
2 Enter the *Name* of `Static_NAT_1to1`.
3 Select the *External Interface* of *port 1* from the drop-down list.
4 Enter the *External IP Address* of `192.168.37.4`.
5 Enter the *Mapped IP Address* range of `10.10.10.42` to `10.10.10.46`.
6 Select OK.

To configure the virtual IP - CLI

```
config firewall vip
  edit "Static_NAT_1to1"
    set extip 192.168.37.4
    set extintf "port1"
    set mappedip 10.10.10.42-10.10.10.46
  next
end
```

Next, configure the firewall policies. Even if no connection needs to be initiated from external to internal, a second security policy number is required to activate the VIP range. Otherwise the IP address of the physical interface is used for NAT. In this example it is set as a “DENY” security policy for security purpose.

To configure the firewall policies - web-based manager

1 Go to *Policy* > *Policy* > *Policy* and select *Create New*.
2 Complete the following and select *OK*:

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT
NAT	Select

3 Complete the following and select *OK*:

Source Interface/Zone	port 1
Source Address	all
Destination Interface/Zone	port 2
Destination Address	Static_NAT_1to1

Schedule	always
Service	ALL
Action	deny
Comments	Used to activate static Source NAT 1-to-1

To configure the firewall policies - CLI

```
config firewall policy
  edit 1
    set srcintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  next
  edit 2
    set srcintf port1
    set dstintf port2
    set srcaddr all"
    set dstaddr Static_NAT_1to1
    set schedule always
    set service ANY
    set action deny
    set comments (Used to activate static Source NAT 1-to-1)
    next
  end
end
```

Traffic shaping and per-IP traffic shaping

Traffic shaping helps to optimize traffic flow through the FortiGate unit, and per-IP traffic shaping does much the same, however, it is applies traffic shaping per IP address instead of per policy or per shaper. Traffic shaping, when included in a security policy, controls the bandwidth available to the policy, and sets the priority of the traffic processed by the policy. Traffic shaping makes it possible to control which policies have the highest priority when large amounts of data are moving through the FortiGate unit. For example, the policy for the corporate web server might be given higher priority than the policies for most employee's computers. An employee who needs extra high speed Internet access could have a special outgoing policy set up with higher bandwidth.

Traffic shaping is available for security policies whose Action is ACCEPT, IPSEC, or SSL VPN. It is also available for all supported services, including H.323, TCP, UDP, ICMP, and ESP.

Traffic shaping is used to improve the quality of bandwidth-intensive and sensitive traffic; it also cannot increase the total amount of bandwidth available. The bandwidth available for traffic set in a traffic shaper is used to control data sessions for traffic in both directions.

For more information about traffic shaping, see the Traffic Shaping chapter in the *FortiOS Handbook*.

Endpoint Security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed. Currently, FortiClient Endpoint Security is available for Microsoft Windows 2000 and later only.

For more information about endpoint security, see the UTM chapter in the *FortiOS Handbook*.

Logging traffic

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13 05:23:47 log_id=4 type=traffic subtype=other pri=notice
vd=root status="start" src="10.41.101.20" srcname="10.41.101.20"
src_port=58115 dst="172.20.120.100" dstname="172.20.120.100"
dst_country="N/A" dst_port=137 tran_ip="N/A" tran_port=0
tran_sip="10.31.101.41" tran_sport=58115 service="137/udp"
proto=17 app_type="N/A" duration=0 rule=1 policyid=1 sent=0 rcvd=0
shaper_drop_sent=0 shaper_drop_rcvd=0 perip_drop=0
src_int="internal" dst_int="wan1" SN=97404 app="N/A" app_cat="N/A"
carrier_ep="N/A"
```

If you want to know more about logging, see the Logging and Reporting chapter in the *FortiOS Handbook*. If you want to know more about traffic log messages, see the *FortiGate Log Message Reference*.

Quality of Service (QoS)

The Quality of Service (QoS) feature is an advanced firewall component that applies bandwidth limits and prioritization to traffic. QoS is the capability of the network to adjust some quality aspects for selected flows within your overall network traffic, and may include such techniques as priority-based queuing and traffic policing.

QoS can be implemented for services that include H.323, TCP, UDP, ICMP, and ESP. QoS uses the following techniques:

Traffic policing	Drops packets that do not conform to bandwidth limitations
Traffic shaping	This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.
Queuing	This transmits packets in order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

QoS can be helpful for organizations that are trying to manage their voice and streaming multi-media traffic, which can rapidly consume bandwidth. Both voice and streaming multi-media are sensitive to latency.

For additional information about QoS, see the Traffic Shaping chapter in the *FortiOS Handbook*.

Identity-based security policies

Identity-based security policies, also known as authentication policies, match traffic that requires a supported authentication protocol to trigger the firewall authentication challenge and successfully authenticate network users. Network users authentication can occur using HTTP, HTTPS, FTP, and Telnet protocols as well as through automatic login using NTLM and FSSO, to bypass user intervention.

Identity-based security policies are usually configured for IPsec or SSL VPN traffic since this type of traffic usually requires authentication from network users.

When configuring identity-based policies, you can use schedules to limit network users authentication sessions. For example, example.com has a schedule policy to use P2P applications between noon and 1:00 pm, and a user authentication timeout of 30 minutes. When a user logs in at 12:15 pm, their authentication time logs them off at 12:45 (30 minutes later). You can configure this type of authentication by using the `schedule-timeout` field in the `config firewall policy` command in the CLI.

Identity-based policy positioning

With identity-based security policies, positioning is extremely important. For a typical security policy, the FortiGate unit matches the source, destination and service of the policy. If matched, it acts on that policy. If not, the FortiGate unit moves to the next policy.

With identity-based policies, once the FortiGate unit matches the source and destination addresses, it processes the identity sub-rules for the user groups and services. That is, it acts on the authentication and completes the remainder of that policy and goes no further in the policy list.

The way identity based policies work is that once src/dest are matched, it will process the identity based sub-rules (for lack of a better term) around the user groups and services. It will never process the rest of your rulebase. For this reason, unique security policies should be placed ***before*** an identity-based policy.

For example, consider the following policies:

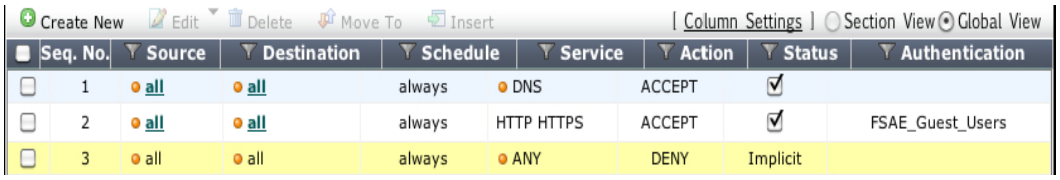

DNS traffic goes through successfully as does any HTTP traffic after being authenticated. However, if there was FTP traffic, it would not get through. As the FortiGate unit processes FTP traffic, it skips rule one since it's matching the source, destination and service. When it moves to rule two it matches the source and destination, it determines there is a match and, sees there are also processes the group/service rules, which requires authentication and acts on those rules. Once satisfied, the FortiGate unit will never go to rule three.

In this situation, where you would want FTP traffic to traverse the FortiGate unit, create a security policy specific to the services you require and place it above the authentication policy.

Identity-based sub-policies

When adding authentication to a security policy, you can add multiple authentication rules, or sub-policies. Within these policies you can include additional UTM profiles, traffic shaping and so on, to take affect on the selected services.

Figure 13: Authentication sub-policies

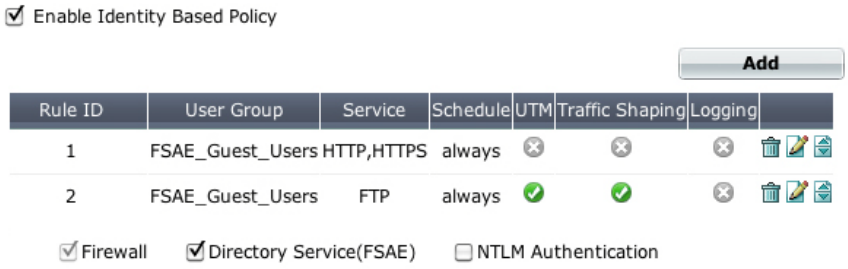

These sub-policies work on the same principle as normal security policies, that is, top down until the criteria has been met. As such, if there is no matching policy within the list, the packet can still be dropped even after authentication is successful.

Appendix

Document conventions

Fortinet technical documentation uses the conventions described below.

IP addresses

To avoid publication of public IP addresses that belong to Fortinet or any other organization, the IP addresses used in Fortinet technical documentation are fictional and follow the documentation guidelines specific to Fortinet. The addresses used are from the private IP address ranges defined in RFC 1918: Address Allocation for Private Internets, available at http://ietf.org/rfc/rfc1918.txt?number-1918.

Most of the examples in this document use the following IP addressing:

- IP addresses are made up of A.B.C.D
- A - can be one of 192, 172, or 10 - the non-public addresses covered in RFC 1918.
- B - 168, or the branch / device / virtual device number.
  - Branch number can be 0xx, 1xx, 2xx - 0 is Head office, 1 is remote, 2 is other.
  - Device or virtual device - allows multiple FortiGate units in this address space (VDOMs).
  - Devices can be from x01 to x99.
- C - interface - FortiGate units can have up to 40 interfaces, potentially more than one on the same subnet
  - 001 - 099- physical address ports, and non -virtual interfaces
  - 100-255 - VLANs, tunnels, aggregate links, redundant links, vdom-links, etc.

- D - usage based addresses, this part is determined by what device is doing
  - The following gives 16 reserved, 140 users, and 100 servers in the subnet.
  - 001 - 009 - reserved for networking hardware, like routers, gateways, etc.
  - 010 - 099 - DHCP range - users
  - 100 - 109 - FortiGate devices - typically only use 100
  - 110 - 199 - servers in general (see later for details)
  - 200 - 249 - static range - users
  - 250 - 255 - reserved (255 is broadcast, 000 not used)
  - The D segment servers can be farther broken down into:
    - 110 - 119 - Email servers
    - 120 - 129 - Web servers
    - 130 - 139 - Syslog servers
    - 140 - 149 - Authentication (RADIUS, LDAP, TACACS+, FSAE, etc)
    - 150 - 159 - VoIP / SIP servers / managers
    - 160 - 169 - FortiAnalyzers
    - 170 - 179 - FortiManagers
    - 180 - 189 - Other Fortinet products (FortiScan, FortiDB, etc.)
    - 190 - 199 - Other non-Fortinet servers (NAS, SQL, DNS, DDNS, etc.)
    - Fortinet products, non-FortiGate, are found from 160 - 189.

The following table shows some examples of how to choose an IP number for a device based on the information given. For internal and dmz, it is assumed in this case there is only one interface being used.

Table 8: Examples of the IP numbering

Location and device	Internal	Dmz	External
Head Office, one FortiGate	10.011.101.100	10.011.201.100	172.20.120.191
Head Office, second FortiGate	10.012.101.100	10.012.201.100	172.20.120.192
Branch Office, one FortiGate	10.021.101.100	10.021.201.100	172.20.120.193
Office 7, one FortiGate with 9 VDOMs	10.079.101.100	10.079.101.100	172.20.120.194
Office 3, one FortiGate, web server	n/a	10.031.201.110	n/a
Bob in accounting on the corporate user network (dhcp) at Head Office, one FortiGate	10.0.11.101.200	n/a	n/a
Router outside the FortiGate	n/a	n/a	172.20.120.195

Example Network configuration

The network configuration shown in Figure 17 or variations on it is used for many of the examples in this document. In this example, the 172.20.120.0 network is equivalent to the Internet. The network consists of a head office and two branch offices.

Figure 17: Example network configuration

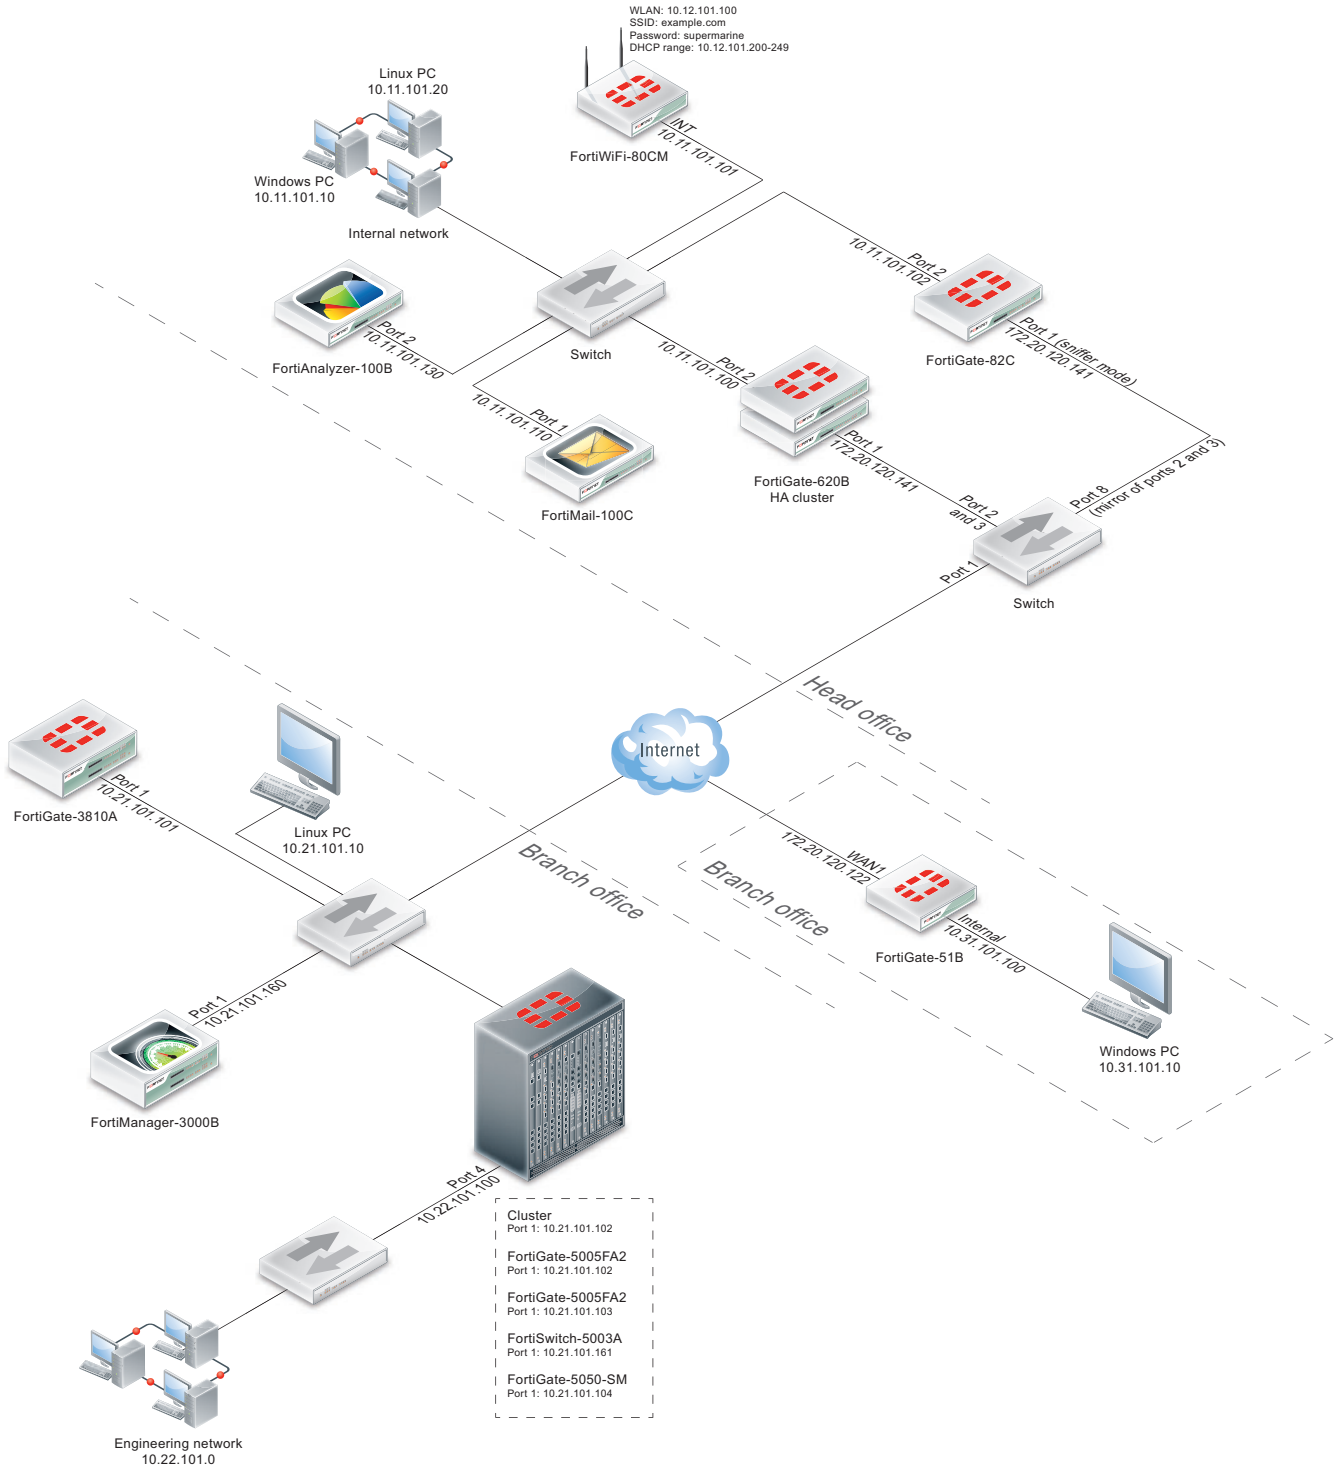

Information highlights

A Must Read item details things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' will not cause data loss but may cause irritation and frustration.

A Troubleshooting tip provides information to help you track down why your configuration is not working.

A Tip provides shortcuts or alternative approaches to the task at hand. Ignoring a tip should have no negative consequences, but you might miss out on a trick that makes your life easier.

Typographical conventions

Table 9: Typographical conventions in Fortinet technical documentation

Convention	Example
Button, menu, text box, field, or check box label	From *Minimum log level*, select *Notification*.

CLI input	`config system dns`
	`  set primary <address_ipv4>`
	`end`

CLI output	`FGT-602803030703 # get system settings`
	`comments  : (null)`
	`opmode  : nat`

Emphasis	HTTP connections are ***not*** secure and can be intercepted by a third party.
File content	`<HTML><HEAD><TITLE>Firewall Authentication</TITLE></HEAD>` `<BODY><H4>You must authenticate to use this service.</H4>`
Hyperlink	Visit the Fortinet Technical Support web site, https://support.fortinet.com.
Keyboard entry	Type a name for the remote VPN peer or client, such as `Central_Office_1`.
Navigation	Go to *VPN* > *IPSEC* > *Auto Key (IKE)*.
Publication	For details, see the *FortiOS Handbook*.

Entering FortiOS configuration data

The configuration of a FortiGate unit is stored as a series of configuration settings in the FortiOS configuration database. To change the configuration you can use the web-based manager or CLI to add, delete or change configuration settings. These configuration changes are stored in the configuration database as you make them.

Individual settings in the configuration database can be text strings, numeric values, selections from a list of allowed options, or on/off (enable/disable).

Entering text strings (names)

Text strings are used to name entities in the configuration. For example, the name of a firewall address, administrative user, and so on. You can enter any character in a FortiGate configuration text string except, to prevent Cross-Site Scripting (XSS) vulnerabilities, text strings in FortiGate configuration names cannot include the following characters:

" (double quote), & (ampersand), ' (single quote), < (less than) and > (greater than)

Most web-based manager text string fields make it easy to add the acceptable amount of text. From the CLI, you can also use the `tree` command to view the number of characters that are allowed. For example, firewall address names can contain up to 64 characters. When you add a firewall address to the web-based manager you are limited to entering 64 characters in the firewall address name field. From the CLI you can do the following to confirm that the firewall address name field allows 64 characters.

```
config firewall address
  tree
  -- [address] --*name (64)
                |- subnet
                |- type
                |- start-ip
                |- end-ip
                |- fqdn (256)
                |- cache-ttl (0,86400)
                |- wildcard
                |- comment (64 xss)
                |- associated-interface (16)
                +- color (0,32)
```

The `tree` command output also shows the number of characters allowed for other firewall address name settings. For example, the fully-qualified domain name (`fqdn`) field can contain up to 256 characters.

Entering numeric values

Numeric values are used to configure various sizes, rates, numeric addresses, or other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or as in the case of MAC or IPv6 addresses separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

Selecting options from a list

If a configuration field can only contain one of a number of selected options, the web-based manager and CLI present you a list of acceptable options and you can select one from the list. No other input is allowed. From the CLI you must spell the selection name correctly.

Enabling or disabling options

If a configuration field can only be on or off (enabled or disabled) the web-based manager presents a check box or other control that can only be enabled or disabled. From the CLI you can set the option to `enable` or `disable`.

Registering your Fortinet product

Access to Fortinet customer services, such as firmware updates, support, and FortiGuard Antivirus and other FortiGuard services, requires product registration.

Take a moment to register your Fortinet product at the Fortinet Technical Support web site, https://support.fortinet.com.

Training

Fortinet Training Services offers courses that orient you quickly to your new equipment, and certifications to verify your knowledge level. Fortinet provides a variety of training programs to serve the needs of our customers and partners world-wide.

Visit the Fortinet Training Services web site at http://campus.training.fortinet.com, or email training@fortinet.com.

Technical Documentation

See the Fortinet Technical Documentation web site, http://docs.fortinet.com, for the most up-to-date technical documentation.

The Fortinet Knowledge Base provides troubleshooting, how-to articles, examples, FAQs, technical notes, and more. Visit the Fortinet Knowledge Base at http://kb.fortinet.com.

Comments on Fortinet technical documentation

Please send information about any errors or omissions in this or any Fortinet technical document to techdoc@fortinet.com.

Customer service and support

Fortinet is committed to your complete satisfaction. Through our regional Technical Assistance Centers and partners worldwide, Fortinet provides remedial support during the operation phase of your Fortinet product's development life cycle. Our Certified Support Partners provide first level technical assistance to Fortinet customers, while the regional TACs solve complex technical issues that our partners are unable to resolve.

Visit the Customer Service and Support web site at http://support.fortinet.com.

Fortinet products End User License Agreement

See the *Fortinet products End User License Agreement*.

Index

A

accept, 44
adding configuring defining
  deny security policy, 50
  how firewalll components create a FortiGate firewall, 12
  how packets flow, 14
  how to apply VLANs and zones to security policies, 23
  how to arrange security poliices, 47
  how to create a basic security policy for Internet access, 52
  interfaces and zones, 23
  ipv4 tunneling, 62
  ipv6, dual stack routing, 62
  remotely connecting to an IPv6 over the Internet, 62
adding, configuring defining
  central NAT table, 66
address
  CIDR format, 24
  FDQN, 31
  geography-based, 28
  groups, 32
  IP pool, 26
  IP range, 25
  IPv6, 61
  matching, IP pool, 27
addresses
  ipv6, 61
addresses, firewall, 24
AFS3, advanced file security encrypted file
  AFS3, 34
AH, predefined service, 34
ANY
  service, 34
AOL
  service, 34
arranging security policies, 47

B

BGP
  service, 34
blocking
  http access by ip, 73
  port 25, 70

C

central NAT, 65
central NAT table
  configuring, 66
certification, 94
CLI syntax conventions, 90
column settings, security policies, 44
comments, documentation, 94
conventions, 87
Cross-Site Scripting
  protection from, 92
custom services, 33
customer service, 94
CVSPSERVER, concurrent versions system proxy server, 34

D

DCE-RPC
  firewall service, 34
default
  password, 9
deny, 44
deny policy, 50
details, security policies, 44
DHCP (Dynamic Host Configuration Protocol)
  service, 34
DHCP6
  service, 34
DNS
  service, 34
  TTL, 31
document conventions
  CLI syntax, 90
documentation, 94
  commenting on, 94
  conventions, 87
  Fortinet, 94
double NAT example, 79
dual stack routing, ipv6, 62

E

ESP
  service, 34

F

FAQ, 94
FDQN, 31
FINGER
  service, 34
firewall
  applying VLANs and zones to security policies, 23
  central NAT table, 66
  how firewall components create a FortiGate firewall, 12
  interfaces and zones, 23
  ipv6, 61
  predefined services, 33
  what is it, 11
  wildcard addresses, 29
firewall address, 24

firewall addresses
ipv6, 61
wildcard address, 29

firewall policies
adding NAT policies to transparent mode, 74

firewall service
AFS3, 34
AH, 34
ANY, 34
AOL, 34
BGP, 34
CVSPSERVER, 34
DCE-RPC, 34
DHCP, 34
DHCP6, 34
DNS, 34
ESP, 34
FINGER, 34
FTP, 34
FTP_GET, 34
FTP_PUT, 34
GOPHER, 34
GRE, 34
H323, 35
HTTP, 35
HTTPS, 35
ICMP_ANY, 35
IKE, 35
IMAP, 35
INFO_ADRESS, 35
INFO_REQUEST, 35
Internet-Locator-Service, 35
IRC, 35
L2TP, 35
LDAP, 35
MGCP, 35
MS-SQL, 35
MYSQL, 35
NetMeeting, 36
NFS, 35
NNTP, 36
NTP, 36
ONC-RPC, 36
OSPF, 36
PC-Anywhere, 36
PING, 36
PING6, 36
POP3, 36
PPTP, 36
QUAKE, 36
RAUDIO, 36
REXEC, 36
RIP, 36
RLOGIN, 36
RSH, 36
RTSP, 37
SAMBA, 37
SCCP, 37
SIP, 37
SIP-MSNmessenger, 37
SMTP, 37
SNMP, 37
SOCKS, 37
SQUID, 37
SSH, 37
SYSLOG, 37
TALK, 37
TCP, 37
TELNET, 37
TFTP, 37
TIMESTAMP, 37
UDP, 38
UUCP, 38
VDOLIVE, 38
viewing predefined list, 33
VNC, 38
WAIS, 38
WINFRAME, 38
WINS, 38
X-WINDOWS, 38
fixed ports, IP pools, 27
FortiGate documentation
commenting on, 94
FortiGate firewall
creating, 12
FortiGuard
Antispam, 9
Antivirus, 9, 93
services, 93
Fortinet
Knowledge Center, 94
Technical Documentation, 94
Technical Documentation, conventions, 87
Technical Support, 94
Technical Support, registering with, 93
Technical Support, web site, 93
Training Services, 94
Fortinet customer service, 94
Fortinet documentation, 94
Fortinet Knowledge Center, 94
FortiOS
ipv6, 61
FTP
service, 34
FTP_GET
service, 34
FTP_PUT
service, 34

G

geography-based addressing, 28
glossary, 94
GOPHER
service, 34
GRE
service, 34
groups, addressing, 32

H

H323
service, 35
how to allow DNS queries to only one DNS server, 50
how to apply VLANs and zones to security policies, 23
how to arrange policies, 47
how to create basic security policy for Internet access, 52
how to test basic security, 53
how to use match-vip, 33
how to use UTM profiles to monitor and protect your network, 40

how-to, 94
HTTP
  service, 35
http
  blocking, 73
HTTPS
  service, 35

I

ICMP processing, 74
ICMP_ANY
  service, 35
identity-based policy, 47
  position, 85
IKE
  service, 35
IMAP
  service, 35
INFO_ADDRESS
  service, 35
INFO_REQUEST
  service, 35
interfaces, 23
  ANY, ANY interface option, 45
Internet-Locator-Service
  service, 35
introduction
  Fortinet documentation, 94
IP address
  private network, 87
IP pool, 26
  address matching, 27
  policies and fixed ports, 27
IP range, 25
IPsec, 44
ipv4 tunneling configuration, ipv6, 62
IPv6, 61
ipv6, 61
  dual stack routing configuration, 62
  ipv4 tunneling configuration, 62
  remotely connecting over the Internet, 62
ipv6 in FortiOS, 61
IRC
  service, 35

K

Knowledge Center, 94

L

L2TP
  service, 35
LDAP
  service, 35
life of a packet, 13
local-in policy, 51

M

M3UA, 68
match-vip, 32
  how to, 33
Message Transfer Part 3, 68
MGCP
  service, 35
mode
  operation, 9
MS-SQL
  service, 35
MTP3 User Adaptation Layer, 68
MYSQL
  service, 35

N

NAT, 65
netmask
  wildcard firewall addresses, 29
NetMeeting
  service, 36
NFS
  service, 35
NNTP
  service, 36
NTP
  service, 36

O

ONC-RPC
  service, 36
operation mode, 9
OSPF
  service, 36

P

packet
  ICMP, 74
  life of, 13
packet flow, 14
packets
  flow, 14
password
  administrator, 9
PC-Anywhere
  service, 36
PING
  service, 36
PING6
  firewall service, 36
policies, 44
  column settings, 44
  expiry, 39
  ICMP packets, 74
  identity-based, 47
  NAT to transparent mode, 74
  order, 45
  timeout, 39
  viewing, 45
policy
  local-in, 51
policy 0, 51

POP3
service, 36
port
blocking port 25, 70
ports
services, 33
position
identity-based policy, 85
PPTP
service, 36
predefined services, 33
product registration, 93
protocol
service, 34
PSTN, 67
Public Switched Telephone Network
See PSTN, 67

Q

QUAKE
service, 36

R

RAUDIO
service, 36
registering
with Fortinet Technical Support, 93
remotely connecting to IPv6 over the Internet, 62
REXEC
firewall service, 36
RFC
1918, 87
RIP
service, 36
RLOGIN
service, 36
RSH
firewall service, 36
RTSP
firewall service, 37

S

SAMBA
service, 37
SCCP
firewall service, 37
schedule
timeout, 39
schedules
expiry, 39
group, 39
one time, 38
recurring, 38
schedule-timeout command, 39
security policies, 51
accept, 44
column settings, 44
deny, 44
deny policy, 50
how to apply VLANs and zones, 23
how to arrange, 47
ICMP packets, 74
identity-based, 47
IPsec, 44
policy order, 45
ssl-vpn policies, 44
viewing, 45
security policy
how to allow Internet access, 52
local-in, 51
verifying traffic is hitting a policy, 53

service
AH, 34
ANY, 34
AOL, 34
BGP, 34
CVSPSERVER, 34
DCE-RPC, 34
DHCP, 34
DHCP6, 34
DNS, 34
ESP, 34
FINGER, 34
FTP, 34
FTP_GET, 34
FTP_PUT, 34
GOPHER, 34
GRE, 34
H323, 35
HTTPS, 35
ICMP_ANY, 35
IKE, 35
IMAP, 35
INFO_ADDRESS, 35
INFO_REQUEST, 35
Internet-Locator-Service, 35
IRC, 35
L2TP, 35
LDAP, 35
MGCP, 35
MS-SQL, 35
MYSQL, 35
NetMeeting, 36
NFS, 35
NNTP, 36
NTP, 36
ONC-RPC, 36
OSPF, 36
PC-Anywhere, 36
PING, 36
PING6, 36
POP3, 36
PPTP, 36
predefined, 33
QUAKE, 36
RAUDIO, 36
REXEC, 36
RIP, 36
RLOGIN, 36
RSH, 36
RTSP, 37
SAMBA, 37
SCCP, 37
service name, 34
SIP, 37
SIP-MSNmessenger, 37
SMTP, 37
SNMP, 37
SOCKS, 37
SQUID, 37
SSH, 37
SYSLOG, 37
TALK, 37
TCP, 37
TELNET, 37
TFTP, 37
TIMESTAMP, 37
UDP, 38
UUCP, 38
VDOLIVE, 38
VNC, 38
WAIS, 38
WINFRAME, 38
WINS, 38
X-WINDOWS, 38
services, 33
custom, 33
list, 33
SIP
service, 37
SIP-MSNmessenger
service, 37
SMTP
service, 37
smtp traffic, 70
SNMP
service, 37
SOCKS
service, 37
SQUID
service, 37
SS7, 68
SSH
service, 37
SSL
service definition, 35, 36
ssl-vpn, 44
SYSLOG
service, 37

T

TALK
service, 37
TCP
service, 37
technical
documentation, 94
documentation conventions, 87
notes, 94
support, 94
technical support, 94
TELNET
service, 37
testing a basic security policy, 53
TFTP
service, 37
TIMESTAMP
service, 37
Training Services, 94
transparent mode
adding NAT policies, 74

U

UDP service, 38
understanding firewall addresses, 24

using UTM profiles to monitor and protect your network, 40
UTM
  profiles, 39
UUCP
  service, 38

V

VDOLIVE
  service, 38
verifying traffic is hitting a policy, 53
viewing
  firewall predefined service list, 33
viewing security policies, 45
vip, 32
vip, grouping, 32
vip, match-vip, 32
virtual ip addresses, 32
VNC
  service, 38
vulnerability
  Cross-Site Scripting, 92
  XSS, 92

W

WAIS
  service, 38
wildcard
  firewall addresses, 29
wildcard addresses, 29
WINFRAME
  service, 38
WINS
  service, 38

X

XSS vulnerability
  protection from, 92
X-WINDOWS
  service, 38

Z

zones, 23